

COMPUTERS & STRUCTURES, INC.

STRUCTURAL AND EARTHQUAKE ENGINEERING SOFTWARE

ETABS® 2016
Integrated Building Design Software

CSI Analysis Reference Manual





CSI Analysis Reference Manual

For SAP2000[®], ETABS[®], SAFE[®]
and CSiBridge[®]

COPYRIGHT

Copyright © Computers & Structures, Inc., 1978-2016
All rights reserved.

The CSI Logo[®], SAP2000[®], ETABS[®], SAFE[®], CSiBridge[®], and SAPFire[®] are registered trademarks of Computers & Structures, Inc. Model-Alive[™] and Watch & Learn[™] are trademarks of Computers & Structures, Inc. Windows[®] is a registered trademark of the Microsoft Corporation. Adobe[®] and Acrobat[®] are registered trademarks of Adobe Systems Incorporated.

The computer programs SAP2000[®], ETABS[®], SAFE[®], and CSiBridge[®] and all associated documentation are proprietary and copyrighted products. Worldwide rights of ownership rest with Computers & Structures, Inc. Unlicensed use of these programs or reproduction of documentation in any form, without prior written authorization from Computers & Structures, Inc., is explicitly prohibited. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior explicit written permission of the publisher.

Further information and copies of this documentation may be obtained from:

Computers & Structures, Inc.
www.csiamerica.com

info@csiamerica.com (for general information)
support@csiamerica.com (for technical support)

DISCLAIMER

CONSIDERABLE TIME, EFFORT AND EXPENSE HAVE GONE INTO THE DEVELOPMENT AND TESTING OF THIS SOFTWARE. HOWEVER, THE USER ACCEPTS AND UNDERSTANDS THAT NO WARRANTY IS EXPRESSED OR IMPLIED BY THE DEVELOPERS OR THE DISTRIBUTORS ON THE ACCURACY OR THE RELIABILITY OF THE PROGRAMS THESE PRODUCTS.

THESE PRODUCTS ARE PRACTICAL AND POWERFUL TOOLS FOR STRUCTURAL DESIGN. HOWEVER, THE USER MUST EXPLICITLY UNDERSTAND THE BASIC ASSUMPTIONS OF THE SOFTWARE MODELING, ANALYSIS, AND DESIGN ALGORITHMS AND COMPENSATE FOR THE ASPECTS THAT ARE NOT ADDRESSED.

THE INFORMATION PRODUCED BY THE SOFTWARE MUST BE CHECKED BY A QUALIFIED AND EXPERIENCED ENGINEER. THE ENGINEER MUST INDEPENDENTLY VERIFY THE RESULTS AND TAKE PROFESSIONAL RESPONSIBILITY FOR THE INFORMATION THAT IS USED.

ACKNOWLEDGMENT

Thanks are due to all of the numerous structural engineers, who over the years have given valuable feedback that has contributed toward the enhancement of this product to its current state.

Special recognition is due Dr. Edward L. Wilson, Professor Emeritus, University of California at Berkeley, who was responsible for the conception and development of the original SAP series of programs and whose continued originality has produced many unique concepts that have been implemented in this version.

Table of Contents

Chapter I	Introduction	I
	Analysis Features	2
	Structural Analysis and Design	3
	About This Manual	3
	Topics	3
	Typographical Conventions	4
	Bold for Definitions	4
	Bold for Variable Data	4
	Italics for Mathematical Variables	4
	Italics for Emphasis	5
	Capitalized Names	5
	Bibliographic References	5
Chapter II	Objects and Elements	7
	Objects	7
	Objects and Elements	8
	Groups	9
Chapter III	Coordinate Systems	II
	Overview	12
	Global Coordinate System	12
	Upward and Horizontal Directions	13
	Defining Coordinate Systems	13
	Vector Cross Product	13
	Defining the Three Axes Using Two Vectors	14

Local Coordinate Systems.	14
Alternate Coordinate Systems.	16
Cylindrical and Spherical Coordinates	17
Chapter IV Joints and Degrees of Freedom	21
Overview	22
Modeling Considerations	23
Local Coordinate System	24
Advanced Local Coordinate System	24
Reference Vectors	25
Defining the Axis Reference Vector	26
Defining the Plane Reference Vector.	26
Determining the Local Axes from the Reference Vectors	27
Joint Coordinate Angles	28
Degrees of Freedom.	30
Available and Unavailable Degrees of Freedom	31
Restrained Degrees of Freedom	32
Constrained Degrees of Freedom.	32
Mixing Restraints and Constraints Not Recommended	32
Active Degrees of Freedom	33
Null Degrees of Freedom.	34
Restraint Supports.	34
Spring Supports	36
Nonlinear Supports	37
Distributed Supports	38
Joint Reactions	39
Base Reactions	39
Masses.	40
Force Load	42
Ground Displacement Load	42
Restraint Displacements	43
Spring Displacements	44
Link/Support Displacements	45
Generalized Displacements	45
Degree of Freedom Output	46
Assembled Joint Mass Output.	47
Displacement Output	47
Force Output	48
Element Joint Force Output	48

Chapter V	Constraints and Welds	49
	Overview	50
	Body Constraint	51
	Joint Connectivity	51
	Local Coordinate System.	51
	Constraint Equations	51
	Plane Definition	52
	Diaphragm Constraint.	53
	Joint Connectivity	53
	Local Coordinate System.	53
	Constraint Equations	54
	Plate Constraint	55
	Joint Connectivity	55
	Local Coordinate System.	55
	Constraint Equations	55
	Axis Definition	56
	Rod Constraint	56
	Joint Connectivity	57
	Local Coordinate System.	57
	Constraint Equations	57
	Beam Constraint.	58
	Joint Connectivity	58
	Local Coordinate System.	59
	Constraint Equations	59
	Equal Constraint.	59
	Joint Connectivity	60
	Local Coordinate System.	60
	Selected Degrees of Freedom	60
	Constraint Equations	60
	Local Constraint.	61
	Joint Connectivity	61
	No Local Coordinate System.	62
	Selected Degrees of Freedom	62
	Constraint Equations	62
	Welds	65
	Automatic Master Joints.	66
	Stiffness, Mass, and Loads	66
	Local Coordinate Systems	67
	Constraint Output	67

Chapter VI	Material Properties	69
	Overview	70
	Local Coordinate System	70
	Stresses and Strains	71
	Isotropic Materials	73
	Uniaxial Materials.	74
	Orthotropic Materials	75
	Anisotropic Materials	75
	Temperature-Dependent Properties	76
	Element Material Temperature	77
	Mass Density	77
	Weight Density	78
	Material Damping	78
	Modal Damping	79
	Viscous Proportional Damping.	80
	Hysteretic Proportional Damping	80
	Nonlinear Material Behavior	80
	Tension and Compression	81
	Shear	81
	Hysteresis.	82
	Application	83
	Friction and Dilatational Angles	84
	Hysteresis Models.	85
	Backbone Curve (Action vs. Deformation)	86
	Cyclic Behavior.	86
	Elastic Hysteresis Model	88
	Kinematic Hysteresis Model	88
	Degrading Hysteresis Model	89
	Takeda Hysteresis Model.	93
	Pivot Hysteresis Model.	94
	Concrete Hysteresis Model.	95
	BRB Hardening Hysteresis Model	97
	Isotropic Hysteresis Model	99
	Modified Darwin-Pecknold Concrete Model	100
	Time-dependent Properties	101
	Properties	101
	Time-Integration Control	102
	Design-Type	102

Chapter VII	The Frame Element	105
	Overview	106
	Joint Connectivity	107
	Insertion Points	107
	Degrees of Freedom	108
	Local Coordinate System	108
	Longitudinal Axis 1	109
	Default Orientation	109
	Coordinate Angle	110
	Advanced Local Coordinate System.	110
	Reference Vector	112
	Determining Transverse Axes 2 and 3	113
	Section Properties	114
	Local Coordinate System	115
	Material Properties	115
	Geometric Properties and Section Stiffnesses.	116
	Shape Type	116
	Automatic Section Property Calculation	118
	Section Property Database Files.	118
	Section-Designer Sections	118
	Additional Mass and Weight	120
	Non-prismatic Sections	120
	Property Modifiers.	123
	Named Property Sets	124
	Insertion Points	125
	Local Axes	126
	End Offsets.	127
	Clear Length.	129
	Rigid-end Factor	129
	Effect upon Non-prismatic Elements	130
	Effect upon Internal Force Output	130
	Effect upon End Releases	130
	End Releases.	131
	Unstable End Releases	132
	Effect of End Offsets	132
	Named Property Sets	132
	Nonlinear Properties	133
	Tension/Compression Limits	133
	Plastic Hinge	134
	Mass	134
	Self-Weight Load	134

Gravity Load	135
Concentrated Span Load	135
Distributed Span Load	137
Loaded Length	137
Load Intensity	137
Projected Loads	137
Temperature Load	140
Strain Load	141
Deformation Load	141
Target-Force Load	142
Internal Force Output	142
Effect of End Offsets	144
Stress Output	144

Chapter VIII Hinge Properties 147

Overview	147
Hinge Properties	149
Hinge Length	150
Plastic Deformation Curve	150
Scaling the Curve	151
Strength Loss	152
Types of P-M2-M3 Hinges	153
Isotropic P-M2-M3 Hinge.	153
Parametric P-M2-M3 Hinge.	156
Fiber P-M2-M3 Hinge	156
Hysteresis Models.	157
Automatic, User-Defined, and Generated Properties	158
Automatic Hinge Properties	159
Analysis Modeling	161
Computational Considerations.	162
Analysis Results	163

Chapter IX The Cable Element 165

Overview	166
Joint Connectivity	167
Undeformed Length	167
Shape Calculator	168
Cable vs. Frame Elements.	169
Number of Segments	170
Degrees of Freedom	170

	Local Coordinate System	170
	Section Properties	171
	Material Properties	171
	Geometric Properties and Section Stiffnesses.	172
	Mass	172
	Self-Weight Load	172
	Gravity Load	173
	Distributed Span Load	173
	Temperature Load	174
	Strain and Deformation Load	174
	Target-Force Load	174
	Nonlinear Analysis.	175
	Element Output	175
Chapter X	The Shell Element	177
	Overview.	178
	Homogeneous	179
	Layered	179
	Joint Connectivity	180
	Shape Guidelines	180
	Edge Constraints	183
	Degrees of Freedom	184
	Local Coordinate System	185
	Normal Axis 3.	186
	Default Orientation	186
	Element Coordinate Angle	186
	Advanced Local Coordinate System.	186
	Reference Vector	188
	Determining Tangential Axes 1 and 2	189
	Section Properties	190
	Area Section Type.	190
	Shell Section Type	190
	Homogeneous Section Properties	191
	Layered Section Property	193
	Property Modifiers.	201
	Named Property Sets	202
	Joint Offsets and Thickness Overwrites	203
	Joint Offsets	203
	Effect of Joint Offsets on the Local Axes.	204
	Thickness Overwrites	205

	Mass	206
	Self-Weight Load	206
	Gravity Load	207
	Uniform Load	207
	Surface Pressure Load	208
	Temperature Load	209
	Strain Load	210
	Internal Force and Stress Output	210
Chapter XI	The Plane Element	215
	Overview	216
	Joint Connectivity	217
	Degrees of Freedom	217
	Local Coordinate System	217
	Stresses and Strains	217
	Section Properties	218
	Section Type	218
	Material Properties	219
	Material Angle	219
	Thickness	219
	Incompatible Bending Modes	220
	Mass	220
	Self-Weight Load	221
	Gravity Load	221
	Surface Pressure Load	222
	Pore Pressure Load	222
	Temperature Load	223
	Stress Output	223
Chapter XII	The Asolid Element	225
	Overview	226
	Joint Connectivity	226
	Degrees of Freedom	227
	Local Coordinate System	227
	Stresses and Strains	228
	Section Properties	228
	Section Type	228
	Material Properties	229
	Material Angle	229

Axis of Symmetry	230
Arc and Thickness.	231
Incompatible Bending Modes.	232
Mass	232
Self-Weight Load	232
Gravity Load.	233
Surface Pressure Load	233
Pore Pressure Load.	234
Temperature Load	234
Rotate Load	234
Stress Output.	235
Chapter XIII The Solid Element	237
Overview.	238
Joint Connectivity	238
Degenerate Solids	239
Degrees of Freedom	240
Local Coordinate System	240
Advanced Local Coordinate System.	240
Reference Vectors.	241
Defining the Axis Reference Vector	241
Defining the Plane Reference Vector	242
Determining the Local Axes from the Reference Vectors	243
Element Coordinate Angles	243
Stresses and Strains	246
Solid Properties	246
Material Properties	246
Material Angles	246
Incompatible Bending Modes.	247
Mass	248
Self-Weight Load	248
Gravity Load.	249
Surface Pressure Load	249
Pore Pressure Load.	249
Temperature Load	250
Stress Output.	250
Chapter XIV The Link/Support Element—Basic	251
Overview.	252

Multi-Linear Plastic Property	289
Hysteretic (Rubber) Isolator Property	291
Friction-Pendulum Isolator Property.	292
Axial Behavior	293
Shear Behavior	294
Linear Behavior.	297
Double-Acting Friction-Pendulum Isolator Property	297
Axial Behavior	297
Shear Behavior	298
Linear Behavior.	299
Triple-Pendulum Isolator Property.	299
Axial Behavior	299
Shear Behavior	300
Linear Behavior.	304
Nonlinear Deformation Loads	304
Frequency-Dependent Link/Support Properties	306
Chapter XVI The Tendon Object	309
Overview.	310
Geometry.	310
Discretization	311
Tendons Modeled as Loads or Elements.	311
Connectivity	311
Degrees of Freedom	312
Local Coordinate Systems	313
Base-line Local Coordinate System.	313
Natural Local Coordinate System.	313
Section Properties	314
Material Properties	314
Geometric Properties and Section Stiffnesses.	314
Tension/Compression Limits	315
Mass	316
Prestress Load	316
Self-Weight Load	317
Gravity Load.	318
Temperature Load	318
Strain Load.	318
Deformation Load	319
Target-Force Load	319
Internal Force Output	320

Chapter XVII	Load Patterns	321
	Overview	322
	Load Patterns, Load Cases, and Load Combinations	323
	Defining Load Patterns	323
	Coordinate Systems and Load Components	324
	Effect upon Large-Displacements Analysis.	324
	Force Load	325
	Ground Displacement Load	325
	Self-Weight Load	325
	Gravity Load	326
	Concentrated Span Load.	327
	Distributed Span Load	327
	Tendon Prestress Load.	327
	Uniform Load	328
	Surface Pressure Load	328
	Pore Pressure Load.	328
	Temperature Load	330
	Strain Load.	331
	Deformation Load	331
	Target-Force Load	331
	Rotate Load	332
	Joint Patterns.	332
	Mass Source	334
	Mass from Specified Load Patterns	335
	Negative Mass.	336
	Multiple Mass Sources	336
	Automated Lateral Loads	338
	Acceleration Loads.	338
Chapter XVIII	Load Cases	341
	Overview	342
	Load Cases	343
	Types of Analysis	343
	Sequence of Analysis	344
	Running Load Cases	345
	Linear and Nonlinear Load Cases	346
	Linear Static Analysis	347
	Multi-Step Static Analysis	348

Linear Buckling Analysis	349
Functions.	350
Load Combinations (Combos).	351
Contributing Cases	351
Types of Combos	352
Examples	353
Correspondence.	354
Additional Considerations.	357
Global Energy Response.	357
Equation Solvers.	361
Environment Variables to Control Analysis.	362
SAPFIRE_NUM_THREADS.	362
SAPFIRE_FILESIZE_MB	363
Accessing the Assembled Stiffness and Mass Matrices	363
Chapter XIX Modal Analysis	365
Overview.	365
Eigenvector Analysis	366
Number of Modes	367
Frequency Range	368
Automatic Shifting	369
Convergence Tolerance.	369
Static-Correction Modes	370
Ritz-Vector Analysis.	372
Number of Modes	373
Starting Load Vectors.	373
Number of Generation Cycles.	375
Modal Analysis Output	375
Periods and Frequencies	376
Participation Factors	376
Participating Mass Ratios	377
Static and Dynamic Load Participation Ratios	378
Chapter XX Response-Spectrum Analysis	383
Overview.	383
Local Coordinate System	385
Response-Spectrum Function	385
Damping.	386
Modal Damping	387
Modal Combination	388

	Periodic and Rigid Response	388
	CQC Method	390
	GMC Method	390
	SRSS Method	390
	Absolute Sum Method	391
	NRC Ten-Percent Method	391
	NRC Double-Sum Method	391
	Directional Combination.	391
	SRSS Method	391
	CQC3 Method.	392
	Absolute Sum Method	393
	Response-Spectrum Analysis Output	394
	Damping and Accelerations.	394
	Modal Amplitudes.	394
	Base Reactions	395
Chapter XXI	Linear Time-History Analysis	397
	Overview.	398
	Loading	398
	Defining the Spatial Load Vectors	399
	Defining the Time Functions	400
	Initial Conditions.	402
	Time Steps	402
	Modal Time-History Analysis	403
	Modal Damping	404
	Direct-Integration Time-History Analysis	405
	Time Integration Parameters	406
	Damping.	406
Chapter XXII	Geometric Nonlinearity	409
	Overview.	409
	Nonlinear Load Cases	411
	The P-Delta Effect	413
	P-Delta Forces in the Frame Element	415
	P-Delta Forces in the Link/Support Element	418
	Other Elements	419
	Initial P-Delta Analysis	419
	Building Structures	420
	Cable Structures.	422
	Guyed Towers.	422
	Large Displacements.	422

Applications	423
Initial Large-Displacement Analysis	423
Chapter XXIII Nonlinear Static Analysis	425
Overview	426
Nonlinearity	426
Important Considerations	427
Loading	428
Load Application Control	428
Load Control	429
Displacement Control	429
Initial Conditions.	430
Output Steps	431
Saving Multiple Steps.	431
Nonlinear Solution Control	433
Maximum Total Steps.	434
Maximum Null (Zero) Steps	434
Event-to-Event Stepping Control	434
Nonlinear Iteration	435
Line Search Option	436
Static Pushover Analysis.	437
Staged Construction	439
Stages	439
Changing Section Properties	442
Output Steps.	442
Example	443
Target-Force Iteration	444
Chapter XXIV Nonlinear Time-History Analysis	447
Overview	448
Nonlinearity	448
Loading	449
Initial Conditions.	449
Time Steps	450
Nonlinear Modal Time-History Analysis (FNA)	451
Initial Conditions	451
Link/Support Effective Stiffness	452
Mode Superposition.	452
Modal Damping.	454
Iterative Solution	455

Static Period	457
Nonlinear Direct-Integration Time-History Analysis	458
Time Integration Parameters	458
Nonlinearity	459
Initial Conditions	459
Damping	459
Nonlinear Solution Control	461

Chapter XXV Frequency-Domain Analyses 465

Overview	466
Harmonic Motion	466
Frequency Domain	467
Damping	468
Sources of Damping	468
Loading	469
Defining the Spatial Load Vectors	470
Frequency Steps	471
Steady-State Analysis	471
Example	472
Power-Spectral-Density Analysis	473
Example	474

Chapter XXVI Moving-Load Analysis 477

Overview for CSiBridge	478
Moving-Load Analysis in SAP2000	479
Bridge Modeler	480
Moving-Load Analysis Procedure	481
Lanes	482
Centerline and Direction	482
Eccentricity	482
Centrifugal Radius	483
Width	483
Interior and Exterior Edges	483
Discretization	484
Influence Lines and Surfaces	485
Vehicle Live Loads	487
Distribution of Loads	487
Axle Loads	487
Uniform Loads	487
Minimum Edge Distances	487

Directions of Loading	488
Restricting a Vehicle to the Lane Length	492
Application of Loads to the Influence Surface	492
Length Effects	494
Application of Loads in Multi-Step Analysis	495
General Vehicle	495
Specification	497
Moving the Vehicle	498
Vehicle Response Components	498
Superstructure (Span) Moment	498
Negative Superstructure (Span) Moment	499
Reactions at Interior Supports	500
Standard Vehicles	500
Vehicle Classes	508
Moving-Load Load Cases	509
Directions of Loading	510
Example 1 — AASHTO HS Loading	512
Example 2 — AASHTO HL Loading	513
Example 3 — Caltrans Permit Loading	514
Example 4 — Restricted Caltrans Permit Loading	516
Example 5 — Eurocode Characteristic Load Model 1	518
Moving Load Response Control	519
Bridge Response Groups	519
Correspondence	520
Influence Line Tolerance	520
Exact and Quick Response Calculation	521
Step-By-Step Analysis	521
Loading	522
Static Analysis	522
Time-History Analysis	523
Enveloping and Load Combinations	524
Computational Considerations	524

Chapter XXVII References

Introduction

SAP2000, ETABS, SAFE, and CSiBridge are software packages from Computers and Structures, Inc. for structural analysis and design. Each package is a fully integrated system for modeling, analyzing, designing, and optimizing structures of a particular type:

- SAP2000 for general structures, including stadiums, towers, industrial plants, offshore structures, piping systems, buildings, dams, soils, machine parts and many others
- ETABS for building structures
- SAFE for floor slabs and base mats
- CSiBridge for bridge structures

At the heart of each of these software packages is a common analysis engine, referred to throughout this manual as SAPfire. This engine is the latest and most powerful version of the well-known SAP series of structural analysis programs. The purpose of this manual is to describe the features of the SAPfire analysis engine.

Throughout this manual reference may be made to the program SAP2000, although it often applies equally to ETABS, SAFE, and CSiBridge. Not all features described will actually be available in every level of each program.

Analysis Features

The SAPfire analysis engine offers the following features:

- Static and dynamic analysis
- Linear and nonlinear analysis
- Dynamic seismic analysis and static pushover analysis
- Vehicle live-load analysis for bridges
- Geometric nonlinearity, including P-delta and large-displacement effects
- Staged (incremental) construction
- Creep, shrinkage, and aging effects
- Buckling analysis
- Steady-state and power-spectral-density analysis
- Frame and shell structural elements, including beam-column, truss, membrane, and plate behavior
- Cable and Tendon elements
- Two-dimensional plane and axisymmetric solid elements
- Three-dimensional solid elements
- Nonlinear link and support elements
- Frequency-dependent link and support properties
- Multiple coordinate systems
- Many types of constraints
- A wide variety of loading options
- Alpha-numeric labels
- Large capacity
- Highly efficient and stable solution algorithms

These features, and many more, make CSI product the state-of-the-art for structural analysis. Note that not all of these features may be available in every level of SAP2000, ETABS, SAFE, and CSiBridge.

Structural Analysis and Design

The following general steps are required to analyze and design a structure using SAP2000, ETABS, SAFE, and CSiBridge:

1. Create or modify a model that numerically defines the geometry, properties, loading, and analysis parameters for the structure
2. Perform an analysis of the model
3. Review the results of the analysis
4. Check and optimize the design of the structure

This is usually an iterative process that may involve several cycles of the above sequence of steps. All of these steps can be performed seamlessly using the SAP2000, ETABS, SAFE, and CSiBridge graphical user interfaces.

About This Manual

This manual describes the theoretical concepts behind the modeling and analysis features offered by the SAPfire analysis engine that underlies the various structural analysis and design software packages from Computers and Structures, Inc. The graphical user interface and the design features are described in separate manuals for each program.

It is imperative that you read this manual and understand the assumptions and procedures used by these software packages before attempting to use the analysis features.

Throughout this manual reference may be made to the program SAP2000, although it often applies equally to ETABS, SAFE, and CSiBridge. Not all features described will actually be available in every level of each program.

Topics

Each Chapter of this manual is divided into topics and subtopics. All Chapters begin with a list of topics covered. These are divided into two groups:

- Basic topics — recommended reading for all users

- Advanced topics — for users with specialized needs, and for all users as they become more familiar with the program.

Following the list of topics is an Overview which provides a summary of the Chapter. Reading the Overview for every Chapter will acquaint you with the full scope of the program.

Typographical Conventions

Throughout this manual the following typographic conventions are used.

Bold for Definitions

Bold roman type (e.g., **example**) is used whenever a new term or concept is defined. For example:

The **global coordinate system** is a three-dimensional, right-handed, rectangular coordinate system.

This sentence begins the definition of the global coordinate system.

Bold for Variable Data

Bold roman type (e.g., **example**) is used to represent variable data items for which you must specify values when defining a structural model and its analysis. For example:

The Frame element coordinate angle, **ang**, is used to define element orientations that are different from the default orientation.

Thus you will need to supply a numeric value for the variable **ang** if it is different from its default value of zero.

Italics for Mathematical Variables

Normal italic type (e.g., *example*) is used for scalar mathematical variables, and bold italic type (e.g., ***example***) is used for vectors and matrices. If a variable data item is used in an equation, bold roman type is used as discussed above. For example:

$$0 \quad \mathbf{da} < \mathbf{db} \quad L$$

Here **da** and **db** are variables that you specify, and L is a length calculated by the program.

Italics for Emphasis

Normal italic type (e.g., *example*) is used to emphasize an important point, or for the title of a book, manual, or journal.

Capitalized Names

Capitalized names (e.g., Example) are used for certain parts of the model and its analysis which have special meaning to SAP2000. Some examples:

- Frame element
- Diaphragm Constraint
- Frame Section
- Load Pattern

Common entities, such as “joint” or “element” are not capitalized.

Bibliographic References

References are indicated throughout this manual by giving the name of the author(s) and the date of publication, using parentheses. For example:

See Wilson and Tetsuji (1983).

It has been demonstrated (Wilson, Yuan, and Dickens, 1982) that ...

All bibliographic references are listed in alphabetical order in Chapter “References” (page 527).

Objects and Elements

The physical structural members in a structural model are represented by objects. Using the graphical user interface, you “draw” the geometry of an object, then “assign” properties and loads to the object to completely define the model of the physical member. For analysis purposes, SAP2000 converts each object into one or more elements.

Basic Topics for All Users

- Objects
- Objects and Elements
- Groups

Objects

The following object types are available, listed in order of geometrical dimension:

- **Point objects**, of two types:
 - **Joint objects:** These are automatically created at the corners or ends of all other types of objects below, and they can be explicitly added to represent supports or to capture other localized behavior.

- **Grounded (one-joint) link/support objects:** Used to model special support behavior such as isolators, dampers, gaps, multi-linear springs, and more.
- **Line objects**, of four types
 - **Frame objects:** Used to model beams, columns, braces, and trusses
 - **Cable objects:** Used to model slender cables under self weight and tension
 - **Tendon objects:** Used to prestressing tendons within other objects
 - **Connecting (two-joint) link/support objects:** Used to model special member behavior such as isolators, dampers, gaps, multi-linear springs, and more. Unlike frame, cable, and tendon objects, connecting link objects can have zero length.
- **Area objects:** Shell elements (plate, membrane, and full-shell) used to model walls, floors, and other thin-walled members; as well as two-dimensional solids (plane-stress, plane-strain, and axisymmetric solids).
- **Solid objects:** Used to model three-dimensional solids.

As a general rule, the geometry of the object should correspond to that of the physical member. This simplifies the visualization of the model and helps with the design process.

Objects and Elements

If you have experience using traditional finite element programs, including earlier versions of SAP2000, ETABS, and SAFE, you are probably used to meshing physical models into smaller finite elements for analysis purposes. Object-based modeling largely eliminates the need for doing this.

For users who are new to finite-element modeling, the object-based concept should seem perfectly natural.

When you run an analysis, SAP2000 automatically converts your object-based model into an element-based model that is used for analysis. This element-based model is called the analysis model, and it consists of traditional finite elements and joints (nodes). Results of the analysis are reported back on the object-based model.

You have control over how the meshing is performed, such as the degree of refinement, and how to handle the connections between intersecting objects. You also have the option to manually mesh the model, resulting in a one-to-one correspondence between objects and elements.

In this manual, the term “element” will be used more often than “object”, since what is described herein is the finite-element analysis portion of the program that operates on the element-based analysis model. However, it should be clear that the properties described here for elements are actually assigned in the interface to the objects, and the conversion to analysis elements is automatic.

One specific case to be aware of is that both one-joint (grounded) link/support objects and two-joint (connecting) link/support objects are always converted into two-joint link/support elements. For the two-joint objects, the conversion to elements is direct. For the one-joint objects, a new joint is created at the same location and is fully restrained. The generated two-joint link/support element is of zero length, with its original joint connected to the structure and the new joint connected to ground by restraints.

Groups

A **group** is a named collection of objects that you define. For each group, you must provide a unique name, then select the objects that are to be part of the group. You can include objects of any type or types in a group. Each object may be part of one of more groups. All objects are always part of the built-in group called “ALL”.

Groups are used for many purposes in the graphical user interface, including selection, design optimization, defining section cuts, controlling output, and more. In this manual, we are primarily interested in the use of groups for defining staged construction. See Topic “Staged Construction” (page 79) in Chapter “Nonlinear Static Analysis” for more information.

Coordinate Systems

Each structure may use many different coordinate systems to describe the location of points and the directions of loads, displacement, internal forces, and stresses. Understanding these different coordinate systems is crucial to being able to properly define the model and interpret the results.

Basic Topics for All Users

- Overview
- Global Coordinate System
- Upward and Horizontal Directions
- Defining Coordinate Systems
- Local Coordinate Systems

Advanced Topics

- Alternate Coordinate Systems
- Cylindrical and Spherical Coordinates

Overview

Coordinate systems are used to locate different parts of the structural model and to define the directions of loads, displacements, internal forces, and stresses.

All coordinate systems in the model are defined with respect to a single global coordinate system. Each part of the model (joint, element, or constraint) has its own local coordinate system. In addition, you may create alternate coordinate systems that are used to define locations and directions.

All coordinate systems are three-dimensional, right-handed, rectangular (Cartesian) systems. Vector cross products are used to define the local and alternate coordinate systems with respect to the global system.

SAP2000 always assumes that Z is the vertical axis, with +Z being upward. The upward direction is used to help define local coordinate systems, although local coordinate systems themselves do not have an upward direction.

The locations of points in a coordinate system may be specified using rectangular or cylindrical coordinates. Likewise, directions in a coordinate system may be specified using rectangular, cylindrical, or spherical coordinate directions at a point.

Global Coordinate System

The **global coordinate system** is a three-dimensional, right-handed, rectangular coordinate system. The three axes, denoted X, Y, and Z, are mutually perpendicular and satisfy the right-hand rule.

Locations in the global coordinate system can be specified using the variables x , y , and z . A vector in the global coordinate system can be specified by giving the locations of two points, a pair of angles, or by specifying a coordinate direction. Coordinate directions are indicated using the values X , Y , and Z . For example, +X defines a vector parallel to and directed along the positive X axis. The sign is required.

All other coordinate systems in the model are ultimately defined with respect to the global coordinate system, either directly or indirectly. Likewise, all joint coordinates are ultimately converted to global X, Y, and Z coordinates, regardless of how they were specified.

Upward and Horizontal Directions

SAP2000 always assumes that Z is the vertical axis, with +Z being upward. Local coordinate systems for joints, elements, and ground-acceleration loading are defined with respect to this upward direction. Self-weight loading always acts downward, in the $-Z$ direction.

The X-Y plane is horizontal. The primary horizontal direction is +X. Angles in the horizontal plane are measured from the positive half of the X axis, with positive angles appearing counterclockwise when you are looking down at the X-Y plane.

If you prefer to work with a different upward direction, you can define an alternate coordinate system for that purpose.

Defining Coordinate Systems

Each coordinate system to be defined must have an origin and a set of three, mutually-perpendicular axes that satisfy the right-hand rule.

The origin is defined by simply specifying three coordinates in the global coordinate system.

The axes are defined as vectors using the concepts of vector algebra. A fundamental knowledge of the **vector cross product** operation is very helpful in clearly understanding how coordinate system axes are defined.

Vector Cross Product

A vector may be defined by two points. It has length, direction, and location in space. For the purposes of defining coordinate axes, only the direction is important. Hence any two vectors that are parallel and have the same sense (i.e., pointing the same way) may be considered to be the same vector.

Any two vectors, V_i and V_j , that are not parallel to each other define a plane that is parallel to them both. The location of this plane is not important here, only its orientation. The cross product of V_i and V_j defines a third vector, V_k , that is perpendicular to them both, and hence normal to the plane. The cross product is written as:

$$V_k = V_i \times V_j$$

The length of V_k is not important here. The side of the V_i - V_j plane to which V_k points is determined by the right-hand rule: The vector V_k points toward you if the acute angle (less than 180°) from V_i to V_j appears counterclockwise.

Thus the sign of the cross product depends upon the order of the operands:

$$V_j \times V_i = -V_i \times V_j$$

Defining the Three Axes Using Two Vectors

A right-handed coordinate system R-S-T can be represented by the three mutually-perpendicular vectors V_r , V_s , and V_t , respectively, that satisfy the relationship:

$$V_t = V_r \times V_s$$

This coordinate system can be defined by specifying two non-parallel vectors:

- An axis reference vector, V_a , that is parallel to axis R
- A plane reference vector, V_p , that is parallel to plane R-S, and points toward the positive-S side of the R axis

The axes are then defined as:

$$V_r = V_a$$

$$V_t = V_r \times V_p$$

$$V_s = V_t \times V_r$$

Note that V_p can be any convenient vector parallel to the R-S plane; it does not have to be parallel to the S axis. This is illustrated in Figure 1 (page 15).

Local Coordinate Systems

Each part (joint, element, or constraint) of the structural model has its own local coordinate system used to define the properties, loads, and response for that part. The axes of the local coordinate systems are denoted 1, 2, and 3. In general, the local coordinate systems may vary from joint to joint, element to element, and constraint to constraint.

There is no preferred upward direction for a local coordinate system. However, the upward +Z direction is used to define the default joint and element local coordinate systems with respect to the global or any alternate coordinate system.

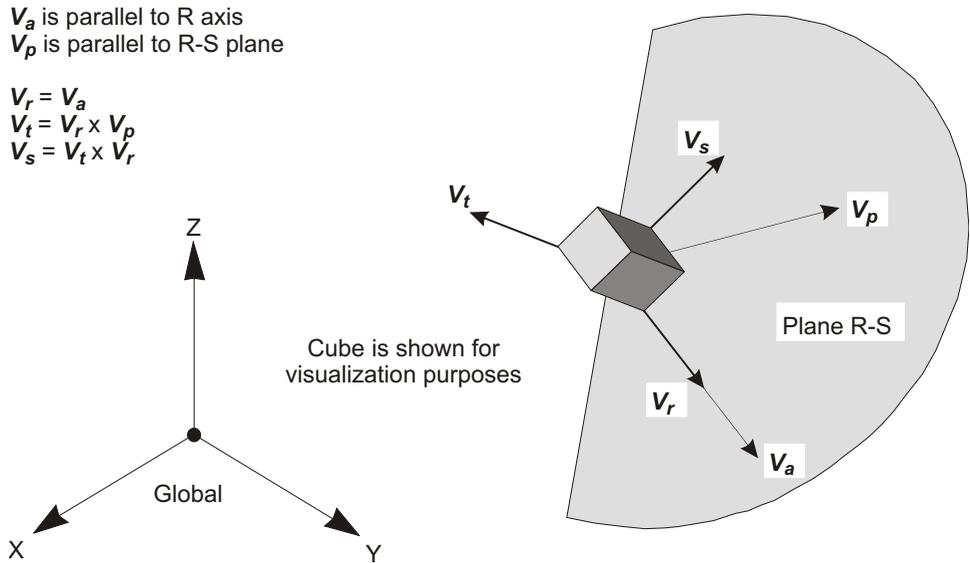


Figure 1

Determining an R-S-T Coordinate System from Reference Vectors V_a and V_p

The joint local 1-2-3 coordinate system is normally the same as the global X-Y-Z coordinate system. However, you may define any arbitrary orientation for a joint local coordinate system by specifying two reference vectors and/or three angles of rotation.

For the Frame, Area (Shell, Plane, and Asolid), and Link/Support elements, one of the element local axes is determined by the geometry of the individual element. You may define the orientation of the remaining two axes by specifying a single reference vector and/or a single angle of rotation. The exception to this is one-joint or zero-length Link/Support elements, which require that you first specify the local-1 (axial) axis.

The Solid element local 1-2-3 coordinate system is normally the same as the global X-Y-Z coordinate system. However, you may define any arbitrary orientation for a solid local coordinate system by specifying two reference vectors and/or three angles of rotation.

The local coordinate system for a Body, Diaphragm, Plate, Beam, or Rod Constraint is normally determined automatically from the geometry or mass distribution of the constraint. Optionally, you may specify one local axis for any Dia-

phragm, Plate, Beam, or Rod Constraint (but not for the Body Constraint); the remaining two axes are determined automatically.

The local coordinate system for an Equal Constraint may be arbitrarily specified; by default it is the global coordinate system. The Local Constraint does not have its own local coordinate system.

For more information:

- See Topic “Local Coordinate System” (page 24) in Chapter “Joints and Degrees of Freedom.”
- See Topic “Local Coordinate System” (page 108) in Chapter “The Frame Element.”
- See Topic “Local Coordinate System” (page 185) in Chapter “The Shell Element.”
- See Topic “Local Coordinate System” (page 217) in Chapter “The Plane Element.”
- See Topic “Local Coordinate System” (page 227) in Chapter “The Asolid Element.”
- See Topic “Local Coordinate System” (page 240) in Chapter “The Solid Element.”
- See Topic “Local Coordinate System” (page 253) in Chapter “The Link/Support Element—Basic.”
- See Chapter “Constraints and Welds (page 49).”

Alternate Coordinate Systems

You may define **alternate coordinate systems** that can be used for locating the joints; for defining local coordinate systems for joints, elements, and constraints; and as a reference for defining other properties and loads. The axes of the alternate coordinate systems are denoted X, Y, and Z.

The global coordinate system and all alternate systems are called **fixed coordinate systems**, since they apply to the whole structural model, not just to individual parts as do the local coordinate systems. Each fixed coordinate system may be used in rectangular, cylindrical or spherical form.

Associated with each fixed coordinate system is a grid system used to locate objects in the graphical user interface. Grids have no meaning in the analysis model.

Each alternate coordinate system is defined by specifying the location of the origin and the orientation of the axes with respect to the global coordinate system. You need:

- The global X, Y, and Z coordinates of the new origin
- The three angles (in degrees) used to rotate from the global coordinate system to the new system

Cylindrical and Spherical Coordinates

The location of points in the global or an alternate coordinate system may be specified using polar coordinates instead of rectangular X-Y-Z coordinates. Polar coordinates include cylindrical CR-CA-CZ coordinates and spherical SB-SA-SR coordinates. See Figure 2 (page 19) for the definition of the polar coordinate systems. Polar coordinate systems are always defined with respect to a rectangular X-Y-Z system.

The coordinates CR, CZ, and SR are lineal and are specified in length units. The coordinates CA, SB, and SA are angular and are specified in degrees.

Locations are specified in *cylindrical* coordinates using the variables **cr**, **ca**, and **cz**. These are related to the rectangular coordinates as:

$$\mathbf{cr} = \sqrt{\mathbf{x}^2 + \mathbf{y}^2}$$

$$\mathbf{ca} = \tan^{-1} \frac{\mathbf{y}}{\mathbf{x}}$$

$$\mathbf{cz} = \mathbf{z}$$

Locations are specified in *spherical* coordinates using the variables **sb**, **sa**, and **sr**. These are related to the rectangular coordinates as:

$$\mathbf{sb} = \tan^{-1} \frac{\sqrt{\mathbf{x}^2 + \mathbf{y}^2}}{\mathbf{z}}$$

$$\mathbf{sa} = \tan^{-1} \frac{\mathbf{y}}{\mathbf{x}}$$

$$\mathbf{sr} = \sqrt{\mathbf{x}^2 + \mathbf{y}^2 + \mathbf{z}^2}$$

A vector in a fixed coordinate system can be specified by giving the locations of two points or by specifying a coordinate direction at a single point P . Coordinate directions are tangential to the coordinate curves at point P . A positive coordinate direction indicates the direction of increasing coordinate value at that point.

Cylindrical coordinate directions are indicated using the values CR , CA , and CZ . Spherical coordinate directions are indicated using the values SB , SA , and SR . The sign is required. See Figure 2 (page 19).

The cylindrical and spherical coordinate directions are not constant but vary with angular position. The coordinate directions do not change with the lineal coordinates. For example, $+SR$ defines a vector directed from the origin to point P .

Note that the coordinates Z and CZ are identical, as are the corresponding coordinate directions. Similarly, the coordinates CA and SA and their corresponding coordinate directions are identical.

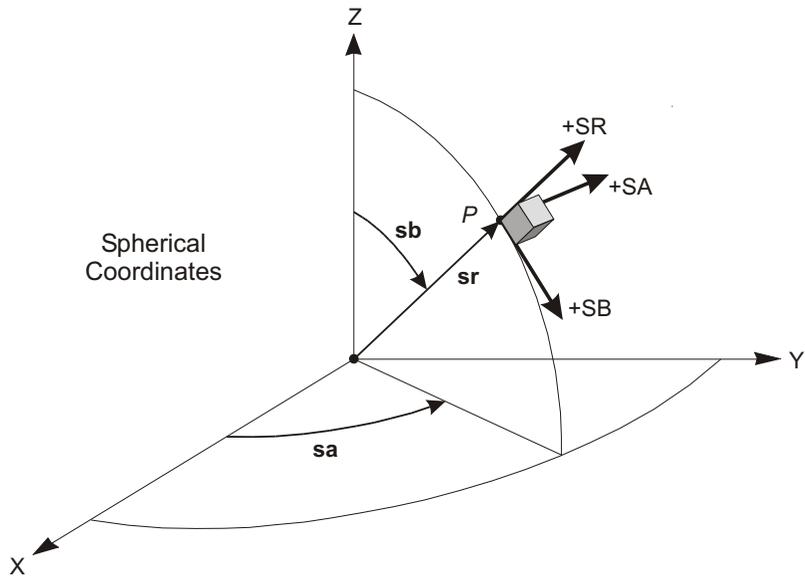
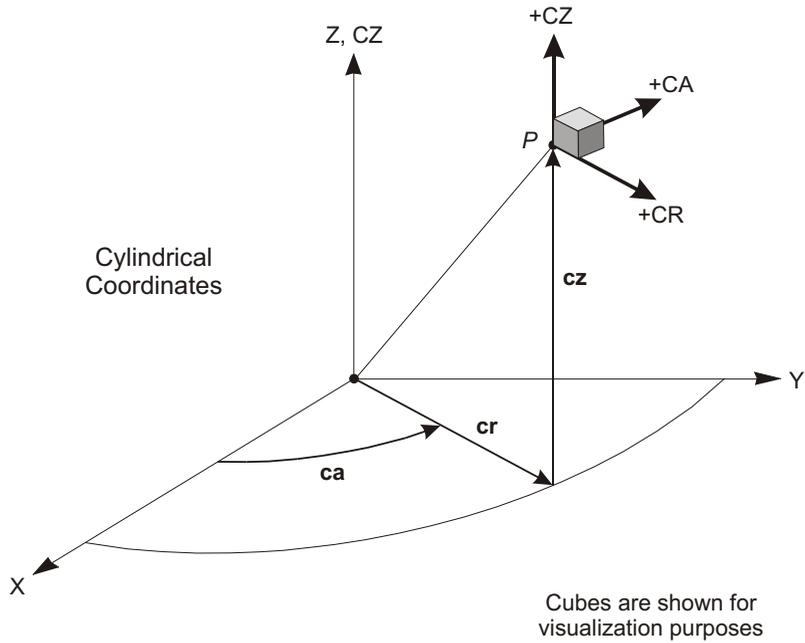


Figure 2
Cylindrical and Spherical Coordinates and Coordinate Directions

Joints and Degrees of Freedom

The **joints** play a fundamental role in the analysis of any structure. Joints are the points of connection between the elements, and they are the primary locations in the structure at which the displacements are known or are to be determined. The displacement components (translations and rotations) at the joints are called the **degrees of freedom**.

This Chapter describes joint properties, degrees of freedom, loads, and output. Additional information about joints and degrees of freedom is given in Chapter “Constraints and Welds” (page 49).

Basic Topics for All Users

- Overview
- Modeling Considerations
- Local Coordinate System
- Degrees of Freedom
- Restraint Supports
- Spring Supports
- Joint Reactions
- Base Reactions

- Masses
- Force Load
- Degree of Freedom Output
- Assembled Joint Mass Output
- Displacement Output
- Force Output

Advanced Topics

- Advanced Local Coordinate System
- Nonlinear Supports
- Distributed Supports
- Ground Displacement Load
- Generalized Displacements
- Element Joint Force Output

Overview

Joints, also known as **nodal points** or **nodes**, are a fundamental part of every structural model. Joints perform a variety of functions:

- All elements are connected to the structure (and hence to each other) at the joints
- The structure is supported at the joints using Restraints and/or Springs
- Rigid-body behavior and symmetry conditions can be specified using Constraints that apply to the joints
- Concentrated loads may be applied at the joints
- Lumped (concentrated) masses and rotational inertia may be placed at the joints
- All loads and masses applied to the elements are actually transferred to the joints
- Joints are the primary locations in the structure at which the displacements are known (the supports) or are to be determined

All of these functions are discussed in this Chapter except for the Constraints, which are described in Chapter “Constraints and Welds” (page 49).

Joints in the analysis model correspond to point objects in the structural-object model. Using the SAP2000, ETABS, SAFE, or CSiBridge graphical user interface, joints (points) are automatically created at the ends of each Line object and at the corners of each Area and Solid object. Joints may also be defined independently of any object.

Automatic meshing of objects will create additional joints corresponding to any elements that are created.

Joints may themselves be considered as elements. Each joint may have its own local coordinate system for defining the degrees of freedom, restraints, joint properties, and loads; and for interpreting joint output. In most cases, however, the global X-Y-Z coordinate system is used as the local coordinate system for all joints in the model. Joints act independently of each other unless connected by other elements.

There are six displacement degrees of freedom at every joint — three translations and three rotations. These displacement components are aligned along the local coordinate system of each joint.

Joints may be loaded directly by concentrated loads or indirectly by ground displacements acting through Restraints, spring supports, or one-joint (grounded) Link/Support objects.

Displacements (translations and rotations) are produced at every joint. Reaction forces and moments acting at each supported joint are also produced.

For more information, see Chapter “Constraints and Welds” (page 49).

Modeling Considerations

The location of the joints and elements is critical in determining the accuracy of the structural model. Some of the factors that you need to consider when defining the elements, and hence the joints, for the structure are:

- The number of elements should be sufficient to describe the geometry of the structure. For straight lines and edges, one element is adequate. For curves and curved surfaces, one element should be used for every arc of 15° or less.
- Element boundaries, and hence joints, should be located at points, lines, and surfaces of discontinuity:
 - Structural boundaries, e.g., corners and edges
 - Changes in material properties

- Changes in thickness and other geometric properties
- Support points (Restrains and Springs)
- Points of application of concentrated loads, except that Frame elements may have concentrated loads applied within their spans
- In regions having large stress gradients, i.e., where the stresses are changing rapidly, an Area- or Solid-element mesh should be refined using small elements and closely-spaced joints. This may require changing the mesh after one or more preliminary analyses.
- More than one element should be used to model the length of any span for which dynamic behavior is important. This is required because the mass is always lumped at the joints, even if it is contributed by the elements.

Local Coordinate System

Each joint has its own **joint local coordinate system** used to define the degrees of freedom, Restraints, properties, and loads at the joint; and for interpreting joint output. The axes of the joint local coordinate system are denoted 1, 2, and 3. By default these axes are identical to the global X, Y, and Z axes, respectively. Both systems are right-handed coordinate systems.

The default local coordinate system is adequate for most situations. However, for certain modeling purposes it may be useful to use different local coordinate systems at some or all of the joints. This is described in the next topic.

For more information:

- See Topic “Upward and Horizontal Directions” (page 13) in Chapter “Coordinate Systems.”
- See Topic “Advanced Local Coordinate System” (page 24) in this Chapter.

Advanced Local Coordinate System

By default, the joint local 1-2-3 coordinate system is identical to the global X-Y-Z coordinate system, as described in the previous topic. However, it may be necessary to use different local coordinate systems at some or all joints in the following cases:

- Skewed Restraints (supports) are present
- Constraints are used to impose rotational symmetry

- Constraints are used to impose symmetry about a plane that is not parallel to a global coordinate plane
- The principal axes for the joint mass (translational or rotational) are not aligned with the global axes
- Joint displacement and force output is desired in another coordinate system

Joint local coordinate systems need only be defined for the affected joints. The global system is used for all joints for which no local coordinate system is explicitly specified.

A variety of methods are available to define a joint local coordinate system. These may be used separately or together. Local coordinate axes may be defined to be parallel to arbitrary coordinate directions in an arbitrary coordinate system or to vectors between pairs of joints. In addition, the joint local coordinate system may be specified by a set of three joint coordinate angles. These methods are described in the subtopics that follow.

For more information:

- See Chapter “Coordinate Systems” (page 11).
- See Topic “Local Coordinate System” (page 24) in this Chapter.

Reference Vectors

To define a joint local coordinate system you must specify two **reference vectors** that are parallel to one of the joint local coordinate planes. The **axis reference vector**, \mathbf{V}_a , must be parallel to one of the local axes ($I = 1, 2, \text{ or } 3$) in this plane and have a positive projection upon that axis. The **plane reference vector**, \mathbf{V}_p , must have a positive projection upon the other local axis ($j = 1, 2, \text{ or } 3, \text{ but } I \neq j$) in this plane, but need not be parallel to that axis. Having a positive projection means that the positive direction of the reference vector must make an angle of less than 90 with the positive direction of the local axis.

Together, the two reference vectors define a local axis, I , and a local plane, i - j . From this, the program can determine the third local axis, k , using vector algebra.

For example, you could choose the axis reference vector parallel to local axis 1 and the plane reference vector parallel to the local 1-2 plane ($I = 1, j = 2$). Alternatively, you could choose the axis reference vector parallel to local axis 3 and the plane reference vector parallel to the local 3-2 plane ($I = 3, j = 2$). You may choose the plane that is most convenient to define using the parameter **local**, which may take on the

values 12, 13, 21, 23, 31, or 32. The two digits correspond to I and j , respectively. The default is value is 31.

Defining the Axis Reference Vector

To define the axis reference vector for joint \mathbf{j} , you must first specify or use the default values for:

- A coordinate direction **axdir** (the default is +Z)
- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system)

You may optionally specify:

- A pair of joints, **axveca** and **axvecb** (the default for each is zero, indicating joint \mathbf{j} itself). If both are zero, this option is not used.

For each joint, the axis reference vector is determined as follows:

1. A vector is found from joint **axveca** to joint **axvecb**. If this vector is of finite length, it is used as the reference vector \mathbf{V}_a
2. Otherwise, the coordinate direction **axdir** is evaluated at joint \mathbf{j} in fixed coordinate system **csys**, and is used as the reference vector \mathbf{V}_a

Defining the Plane Reference Vector

To define the plane reference vector for joint \mathbf{j} , you must first specify or use the default values for:

- A primary coordinate direction **pldirp** (the default is +X)
- A secondary coordinate direction **pldirs** (the default is +Y). Directions **pldirs** and **pldirp** should not be parallel to each other unless you are sure that they are not parallel to local axis 1
- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system). This will be the same coordinate system that was used to define the axis reference vector, as described above

You may optionally specify:

- A pair of joints, **plveca** and **plvecb** (the default for each is zero, indicating joint \mathbf{j} itself). If both are zero, this option is not used.

For each joint, the plane reference vector is determined as follows:

1. A vector is found from joint **plveca** to joint **plvecb**. If this vector is of finite length and is not parallel to local axis I , it is used as the reference vector \mathbf{V}_p
2. Otherwise, the primary coordinate direction **pldirp** is evaluated at joint **j** in fixed coordinate system **csys**. If this direction is not parallel to local axis I , it is used as the reference vector \mathbf{V}_p
3. Otherwise, the secondary coordinate direction **pldirs** is evaluated at joint **j** in fixed coordinate system **csys**. If this direction is not parallel to local axis I , it is used as the reference vector \mathbf{V}_p
4. Otherwise, the method fails and the analysis terminates. This will never happen if **pldirp** is not parallel to **pldirs**

A vector is considered to be parallel to local axis I if the sine of the angle between them is less than 10^{-3} .

Determining the Local Axes from the Reference Vectors

The program uses vector cross products to determine the local axes from the reference vectors. The three axes are represented by the three unit vectors \mathbf{V}_1 , \mathbf{V}_2 and \mathbf{V}_3 , respectively. The vectors satisfy the cross-product relationship:

$$\mathbf{V}_1 \times \mathbf{V}_2 = \mathbf{V}_3$$

The local axis \mathbf{V}_i is given by the vector \mathbf{V}_a after it has been normalized to unit length.

The remaining two axes, \mathbf{V}_j and \mathbf{V}_k , are defined as follows:

- If I and j permute in a positive sense, i.e., **local** = 12, 23, or 31, then:

$$\mathbf{V}_k = \mathbf{V}_i \times \mathbf{V}_p \quad \text{and} \\ \mathbf{V}_j = \mathbf{V}_k \times \mathbf{V}_i$$

- If I and j permute in a negative sense, i.e., **local** = 21, 32, or 13, then:

$$\mathbf{V}_k = \mathbf{V}_p \times \mathbf{V}_i \quad \text{and} \\ \mathbf{V}_j = \mathbf{V}_i \times \mathbf{V}_k$$

An example showing the determination of the joint local coordinate system using reference vectors is given in Figure 3 (page 28).

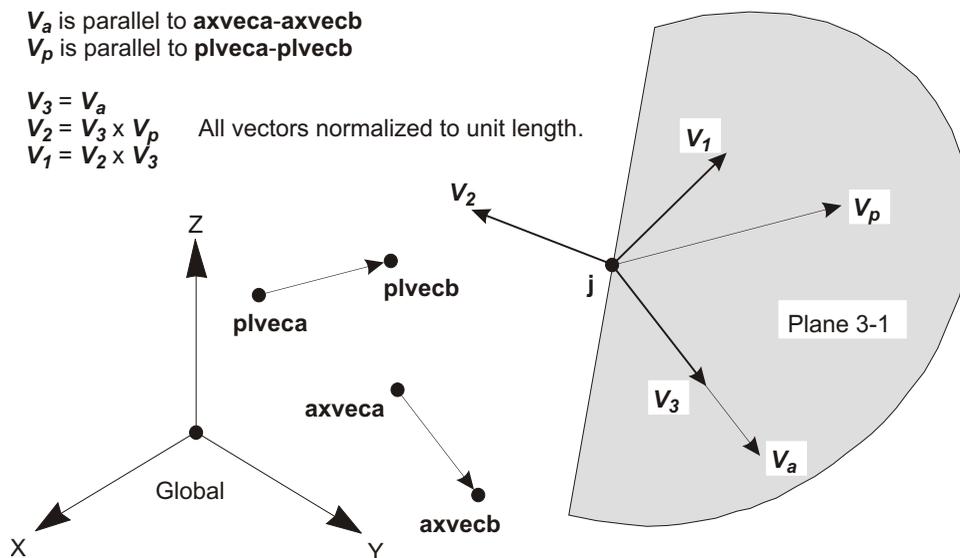


Figure 3
*Example of the Determination of the Joint Local Coordinate System
 Using Reference Vectors for **local=31***

Joint Coordinate Angles

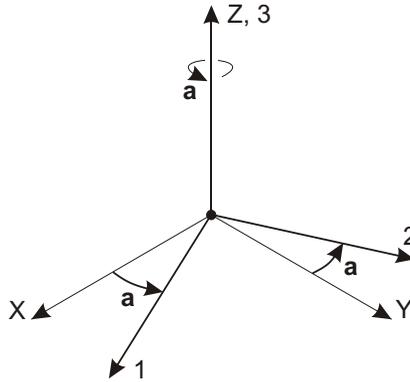
The joint local coordinate axes determined from the reference vectors may be further modified by the use of three **joint coordinate angles**, denoted **a**, **b**, and **c**. In the case where the default reference vectors are used, the joint coordinate angles define the orientation of the joint local coordinate system with respect to the global axes.

The joint coordinate angles specify rotations of the local coordinate system about its own current axes. The resulting orientation of the joint local coordinate system is obtained according to the following procedure:

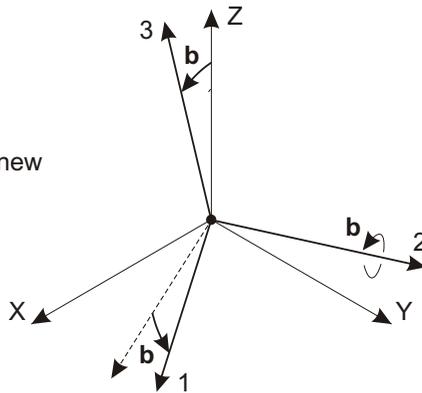
1. The local system is first rotated about its +3 axis by angle **a**
2. The local system is next rotated about its resulting +2 axis by angle **b**
3. The local system is lastly rotated about its resulting +1 axis by angle **c**

The order in which the rotations are performed is important. The use of coordinate angles to orient the joint local coordinate system with respect to the global system is shown in Figure 4 (page 29).

Step 1: Rotation about local 3 axis by angle a



Step 2: Rotation about new local 2 axis by angle b



Step 3: Rotation about new local 1 axis by angle c

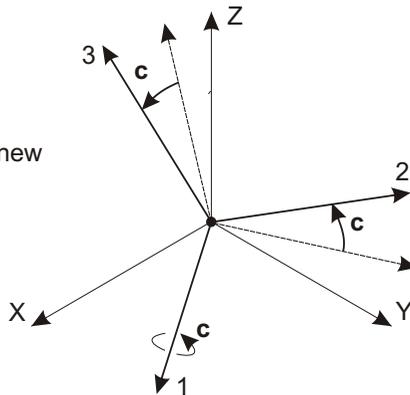


Figure 4

Use of Joint Coordinate Angles to Orient the Joint Local Coordinate System

Degrees of Freedom

The deflection of the structural model is governed by the displacements of the joints. Every joint of the structural model may have up to six displacement components:

- The joint may translate along its three local axes. These translations are denoted U1, U2, and U3.
- The joint may rotate about its three local axes. These rotations are denoted R1, R2, and R3.

These six displacement components are known as the **degrees of freedom** of the joint. In the usual case where the joint local coordinate system is parallel to the global system, the degrees of freedom may also be identified as UX, UY, UZ, RX, RY and RZ, according to which global axes are parallel to which local axes. The joint local degrees of freedom are illustrated in Figure 5 (page 31).

In addition to the regular joints that you explicitly define as part of your structural model, the program automatically creates master joints that govern the behavior of any Constraints and Welds that you may have defined. Each master joint has the same six degrees of freedom as do the regular joints. See Chapter “Constraints and Welds” (page 49) for more information.

Each degree of freedom in the structural model must be one of the following types:

- Active — the displacement is computed during the analysis
- Restrained — the displacement is specified, and the corresponding reaction is computed during the analysis
- Constrained — the displacement is determined from the displacements at other degrees of freedom
- Null — the displacement does not affect the structure and is ignored by the analysis
- Unavailable — the displacement has been explicitly excluded from the analysis

These different types of degrees of freedom are described in the following subtopics.

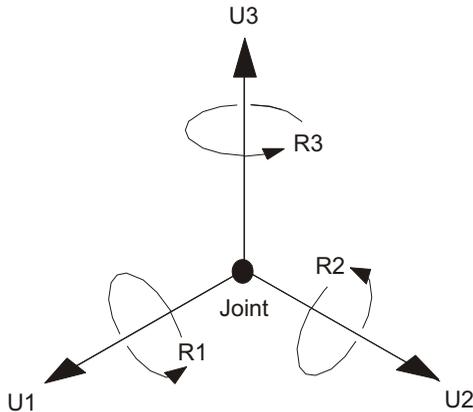


Figure 5

The Six Displacement Degrees of Freedom in the Joint Local Coordinate System

Available and Unavailable Degrees of Freedom

You may explicitly specify the global degrees of freedom that are available to every joint in the structural model. By default, all six degrees of freedom are available to every joint. This default should generally be used for all three-dimensional structures.

For certain planar structures, however, you may wish to restrict the **available degrees of freedom**. For example, in the X-Y plane: a planar truss needs only UX and UY; a planar frame needs only UX, UY, and RZ; and a planar grid or flat plate needs only UZ, RX, and RY.

The degrees of freedom that are not specified as being available are called **unavailable degrees of freedom**. Any stiffness, loads, mass, Restraints, or Constraints that are applied to the unavailable degrees of freedom are ignored by the analysis.

The available degrees of freedom are always referred to the global coordinate system, and they are the same for every joint in the model. If any joint local coordinate systems are used, they must not couple available degrees of freedom with the unavailable degrees of freedom at any joint. For example, if the available degrees of freedom are UX, UY, and RZ, then all joint local coordinate systems must have one local axis parallel to the global Z axis.

Restrained Degrees of Freedom

If the displacement of a joint along any one of its available degrees of freedom is known, such as at a support point, that degree of freedom is **restrained**. The known value of the displacement may be zero or non-zero, and may be different in different Load Patterns. The force along the restrained degree of freedom that is required to impose the specified restraint displacement is called the reaction, and is determined by the analysis.

Unavailable degrees of freedom are essentially restrained. However, they are excluded from the analysis and no reactions are computed, even if they are non-zero.

See Topic “Restraint Supports” (page 34) in this Chapter for more information.

Constrained Degrees of Freedom

Any joint that is part of a Constraint or Weld may have one or more of its available degrees of freedom **constrained**. The program automatically creates a master joint to govern the behavior of each Constraint, and a master joint to govern the behavior of each set of joints that are connected together by a Weld. The displacement of a constrained degree of freedom is then computed as a linear combination of the displacements along the degrees of freedom at the corresponding master joint.

If a constrained degree of freedom is also restrained, the restraint will be applied to the constraint as a whole.

See Chapter “Constraints and Welds” (page 49) for more information.

Mixing Restraints and Constraints Not Recommended

It is *not* recommended that restrained degrees of freedom also be constrained, although it is permitted. Reactions computed at such degrees of freedom will not include the contributions to the reaction from joints connected by constraints. Ground displacement loads applied at the constraint will not be applied to the joints connected by constraint. For better results, use springs or grounded (one-joint) link/support objects to support joints that are also constrained.

Similarly, it is *not* recommended, that a given degree of freedom be included in more than one constraint, although it is permitted. The analysis will attempt to combine the coupled constraints, but the results may not be as accurate as using a single constraint for all coupled joints, especially for dynamics. For better results, include all coupled joints in a single constraint whenever possible.

Note that using fixed degrees of freedom in a link/support property is the same as specifying a constraint. For this reason, such link/support objects should not be connected together or connected to constrained joints. In such cases, it is better to use large (but not too large) stiffnesses rather than fixed degrees of freedom in the link/support property definition.

Active Degrees of Freedom

All available degrees of freedom that are neither constrained nor restrained must be either active or null. The program will automatically determine the **active degrees of freedom** as follows:

- If any load or stiffness is applied along *any* translational degree of freedom at a joint, then all available translational degrees of freedom at that joint are made active unless they are constrained or restrained.
- If any load or stiffness is applied along *any* rotational degree of freedom at a joint, then all available rotational degrees of freedom at that joint are made active unless they are constrained or restrained.
- All degrees of freedom at a master joint that govern constrained degrees of freedom are made active.

A joint that is connected to any element or to a translational spring will have all of its translational degrees of freedom activated. A joint that is connected to a Frame, Shell, or Link/Support element, or to any rotational spring will have all of its rotational degrees of freedom activated. An exception is a Frame element with only truss-type stiffness, which will not activate rotational degrees of freedom.

Every active degree of freedom has an associated equation to be solved. If there are N active degrees of freedom in the structure, there are N equations in the system, and the structural stiffness matrix is said to be of order N . The amount of computational effort required to perform the analysis increases with N .

The load acting along each active degree of freedom is known (it may be zero). The corresponding displacement will be determined by the analysis.

If there are active degrees of freedom in the system at which the stiffness is known to be zero, such as the out-of-plane translation in a planar-frame, these must either be restrained or made unavailable. Otherwise, the structure is unstable and the solution of the static equations will fail.

For more information:

- See Topic “Springs” (page 36) in this Chapter.

- See Topic “Degrees of Freedom” (page 108) in Chapter “The Frame Element.”
- See Topic “Degrees of Freedom” (page 170) in Chapter “The Cable Element.”
- See Topic “Degrees of Freedom” (page 181) in Chapter “The Shell Element.”
- See Topic “Degrees of Freedom” (page 217) in Chapter “The Plane Element.”
- See Topic “Degrees of Freedom” (page 227) in Chapter “The Asolid Element.”
- See Topic “Degrees of Freedom” (page 240) in Chapter “The Solid Element.”
- See Topic “Degrees of Freedom” (page 253) in Chapter “The Link/Support Element—Basic.”
- See Topic “Degrees of Freedom” (page 312) in Chapter “The Tendon Object.”

Null Degrees of Freedom

The available degrees of freedom that are not restrained, constrained, or active, are called the null degrees of freedom. Because they have no load or stiffness, their displacements and reactions are zero, and they have no effect on the rest of the structure. The program automatically excludes them from the analysis.

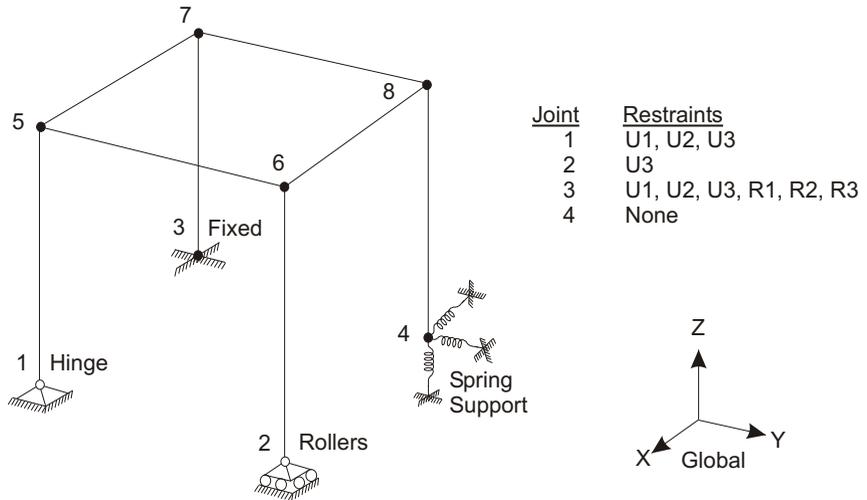
Joints that have no elements connected to them typically have all six degrees of freedom null. Joints that have only solid-type elements (Plane, Asolid, and Solid) connected to them typically have the three rotational degrees of freedom null.

Restraint Supports

If the displacement of a joint along any of its available degrees of freedom has a known value, either zero (e.g., at support points) or non-zero (e.g., due to support settlement), a **Restraint** must be applied to that degree of freedom. The known value of the displacement may differ from one Load Pattern to the next, but the degree of freedom is restrained for all Load Patterns. In other words, it is not possible to have the displacement known in one Load Pattern and unknown (unrestrained) in another Load Pattern.

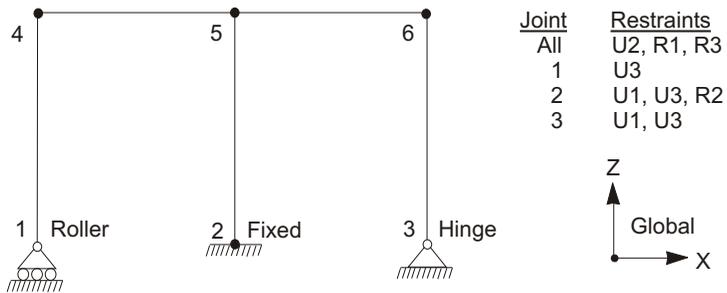
Restraints should also be applied to any available degrees of freedom in the system at which the stiffness is known to be zero, such as the out-of-plane translation and in-plane rotations of a planar-frame. Otherwise, the structure is unstable and the solution of the static equations will complain.

Restraints are always applied to the joint local degrees of freedom U1, U2, U3, R1, R2, and R3. Examples of Restraints are shown in Figure 6 (page 35).



3-D Frame Structure

Notes: Joints are indicated with dots:
 • Solid dots indicate moment continuity
 ○ Open dots indicate hinges
 All joint local 1-2-3 coordinate systems are identical to the global X-Y-Z coordinate system



2-D Frame Structure, X-Z plane

Figure 6
Examples of Restraints

If a restraint is applied to an unavailable degree of freedom, it is ignored. The displacement will be zero, but no reaction will be computed.

In general, you should not apply restraints to constrained degrees of freedom. However, if you do, the analysis will attempt to automatically rewrite the constraint equations to accommodate the restraint. Reactions computed at such degrees of freedom will not include the contributions to the reaction from joints connected by constraints. Ground displacement loads applied at the constraint will not be applied to the joints connected by constraint. For better results, use springs or grounded (one-joint) Link/Supports to support joints that are also constrained.

For more information:

- See Topic “Degrees of Freedom” (page 30) in this Chapter.
- See Topic “Restraint Displacement Load” (page 42) in this Chapter.

Spring Supports

Any of the six degrees of freedom at any of the joints in the structure can have translational or rotational spring support conditions. These springs elastically connect the joint to the ground. Spring supports along restrained degrees of freedom do not contribute to the stiffness of the structure.

Springs may be specified that couple the degrees of freedom at a joint. The spring forces that act on a joint are related to the displacements of that joint by a 6x6 symmetric matrix of spring stiffness coefficients. These forces tend to oppose the displacements.

Spring stiffness coefficients may be specified in the global coordinate system, an Alternate Coordinate System, or the joint local coordinate system.

In a joint local coordinate system, the spring forces and moments F_1 , F_2 , F_3 , M_1 , M_2 and M_3 at a joint are given by:

$$\begin{array}{rcccccccc}
 F_1 & \mathbf{u1} & \mathbf{u1u2} & \mathbf{u1u3} & \mathbf{u1r1} & \mathbf{u1r2} & \mathbf{u1r3} & u_1 & \\
 F_2 & & \mathbf{u2} & \mathbf{u2u3} & \mathbf{u2r1} & \mathbf{u2r2} & \mathbf{u2r3} & u_2 & \\
 F_3 & & & \mathbf{u3} & \mathbf{u3r1} & \mathbf{u3r2} & \mathbf{u3r3} & u_3 & \\
 M_1 & & & & \mathbf{r1} & \mathbf{r1r2} & \mathbf{r1r3} & r_1 & \\
 M_2 & & \mathbf{sym.} & & & \mathbf{r2} & \mathbf{r2r3} & r_2 & \\
 M_3 & & & & & & \mathbf{r3} & r_3 &
 \end{array} \quad (\text{Eqn. 1})$$

where u_1, u_2, u_3, r_1, r_2 and r_3 are the joint displacements and rotations, and the terms $\mathbf{u1}, \mathbf{u1u2}, \mathbf{u2}, \dots$ are the specified spring stiffness coefficients.

In any fixed coordinate system, the spring forces and moments F_x, F_y, F_z, M_x, M_y and M_z at a joint are given by:

$$\begin{array}{cccccccc}
 F_x & \mathbf{ux} & \mathbf{uxuy} & \mathbf{uxuz} & \mathbf{uxrx} & \mathbf{uxry} & \mathbf{uxrz} & u_x \\
 F_y & & \mathbf{uy} & \mathbf{uyuz} & \mathbf{uyrx} & \mathbf{uyry} & \mathbf{uyrz} & u_y \\
 F_z & & & \mathbf{uz} & \mathbf{uzrx} & \mathbf{uzry} & \mathbf{uzrz} & u_z \\
 M_x & & & & \mathbf{rx} & \mathbf{rxry} & \mathbf{rxrz} & r_x \\
 M_y & & \mathbf{sym.} & & & \mathbf{ry} & \mathbf{ryrz} & r_y \\
 M_z & & & & & & \mathbf{rz} & r_z
 \end{array}$$

where u_x, u_y, u_z, r_x, r_y and r_z are the joint displacements and rotations, and the terms $\mathbf{ux}, \mathbf{uxuy}, \mathbf{uy}, \dots$ are the specified spring stiffness coefficients.

For springs that do not couple the degrees of freedom in a particular coordinate system, only the six diagonal terms need to be specified since the off-diagonal terms are all zero. When coupling is present, all 21 coefficients in the upper triangle of the matrix must be given; the other 15 terms are then known by symmetry.

If the springs at a joint are specified in more than one coordinate system, standard coordinate transformation techniques are used to convert the 6x6 spring stiffness matrices to the joint local coordinate system, and the resulting stiffness matrices are then added together on a term-by-term basis. The final spring stiffness matrix at each joint in the structure should have a determinant that is zero or positive. Otherwise the springs may cause the structure to be unstable.

The displacement of the grounded end of the spring may be specified to be zero or non-zero (e.g., due to support settlement). This **spring displacement** may vary from one Load Pattern to the next.

For more information:

- See Topic “Degrees of Freedom” (page 30) in this Chapter.
- See Topic “Spring Displacement Load” (page 43) in this Chapter.

Nonlinear Supports

In certain versions of the program, you may define nonlinear supports at the joints using the Link/Support element. Nonlinear support conditions that can be modeled

include gaps (compression only), multi-linear elastic or plastic springs, viscous dampers, base isolators, and more.

This Link/Support can be used in two ways:

- You can add (draw) a one-joint object, in which case it is considered a Support object, and it connects the joint directly to the ground.
- The object can also be drawn with two joints, in which case it is considered a Link object. You can use a Link object as a support if you connect one end to the structure, and fully restrain the other end.

Both methods have the same effect. During analysis, one-joint Support objects are converted to two-joint Link elements of zero length, and connected to a generated joint that is fully restrained.

Multiple Link/Support elements can be connected to a single joint, in which case they act in parallel. Each Link/Support element has its own element local coordinate system that is independent of the joint local coordinate system.

Restraints and springs may also exist at the joint. Of course, any degree of freedom that is restrained will prevent deformation in the Link/Support element in that direction.

See Chapters “The Link/Support Element – Basic” (page 251) and “The Link/Support Element – Advanced” (page 275) for more information.

Distributed Supports

You may assign distributed spring supports along the length of a Frame element, or over the any face of an area object (Shell, Plane, Asolid) or Solid element. These springs may be linear, multi-linear elastic, or multi-linear plastic. These springs are converted to equivalent two-joint Link/Support elements acting at the joints of the element, after accounting for the tributary length or area of the element. The generated Link/Support elements are of zero length, with one end connected to the parent object, and the other end connected to a generated joint that is fully restrained.

Because these springs act at the joints, it may be necessary to mesh the elements to capture localized effects of such distributed supports. The best way to do this is usually to use the automatic internal meshing options available in the graphical user interface. This allows you to change the meshing easily, while still being able to work with large, simpler model objects.

It is not possible to assign distributed restraint supports directly. However, when using automatic internal meshing, you may optionally specify that the meshed elements use the same restraint conditions that are present on the parent object.

For more information, see Topics “Restraint Supports” (page 34), “Spring Support” (page 36), “Nonlinear Supports” (page 37) in this Chapter, and also Chapter “Objects and Elements” (page 7.)

Joint Reactions

The force or moment along the degree of freedom that is required to enforce any support condition is called the **reaction**, and it is determined by the analysis. The reaction includes the forces (or moments) from all supports at the joint, including restraints, springs, and one-joint Link/Support objects. The tributary effect of any distributed supports is included in the reaction.

If a one-joint Link/Support object is used, the reaction will be reported at the original joint connected to the structure, not at the restrained end of the generated two-joint Link/Support element. The reaction at the generated joint will be reported as zero since it has been transferred to the original joint.

For more information, see Topics “Restraint Supports” (page 34), “Spring Support” (page 36), “Nonlinear Supports” (page 37), and “Distributed Supports” (page 38) in this Chapter.

Base Reactions

Base Reactions are the resultant force and moment of all the joint reactions acting on the structure, computed at the global origin or at some other location that you choose. This produces three force components and three moment components. The base forces are not affected by the chosen location, but the base moments are. For seismic analysis the horizontal forces are called the base shears, and the moments about the horizontal axes are called the overturning moments.

Base reactions are available for all Load Cases and Combos except for Moving-Load Load Cases. The centroids (center of action) are also available for each force component of the base reactions. Note that these are the centroids of the reactions, which may not always be the same as the centroids of the applied load causing the reaction.

For more information, see Topic “Joint Reactions” (page 39) in this Chapter.

Masses

In a dynamic analysis, the mass of the structure is used to compute inertial forces. Normally, the mass is obtained from the elements using the mass density of the material and the volume of the element. This automatically produces lumped (uncoupled) masses at the joints. The element mass values are equal for each of the three translational degrees of freedom. No mass moments of inertia are produced for the rotational degrees of freedom. This approach is adequate for most analyses.

It is often necessary to place additional concentrated masses and/or mass moments of inertia at the joints. These can be applied to any of the six degrees of freedom at any of the joints in the structure.

For computational efficiency and solution accuracy, SAP2000 always uses lumped masses. This means that there is no mass coupling between degrees of freedom at a joint or between different joints. These uncoupled masses are always referred to the local coordinate system of each joint. Mass values along restrained degrees of freedom are ignored.

Inertial forces acting on the joints are related to the accelerations at the joints by a 6x6 matrix of mass values. These forces tend to oppose the accelerations. In a joint local coordinate system, the inertia forces and moments F_1 , F_2 , F_3 , M_1 , M_2 and M_3 at a joint are given by:

$$\begin{array}{rccccccc}
 F_1 & \mathbf{u1} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \ddot{u}_1 \\
 F_2 & & \mathbf{u2} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \ddot{u}_2 \\
 F_3 & & & \mathbf{u3} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \ddot{u}_3 \\
 M_1 & & & & \mathbf{r1} & \mathbf{0} & \mathbf{0} & \ddot{r}_1 \\
 M_2 & & \mathbf{sym.} & & & \mathbf{r2} & \mathbf{0} & \ddot{r}_2 \\
 M_3 & & & & & & \mathbf{r3} & \ddot{r}_3
 \end{array}$$

where \ddot{u}_1 , \ddot{u}_2 , \ddot{u}_3 , \ddot{r}_1 , \ddot{r}_2 and \ddot{r}_3 are the translational and rotational accelerations at the joint, and the terms $\mathbf{u1}$, $\mathbf{u2}$, $\mathbf{u3}$, $\mathbf{r1}$, $\mathbf{r2}$, and $\mathbf{r3}$ are the specified mass values.

Uncoupled joint masses may instead be specified in the global coordinate system, in which case they are transformed to the joint local coordinate system. Coupling terms will be generated during this transformation in the following situation:

- The joint local coordinate system directions are not parallel to global coordinate directions, and
- The three translational masses or the three rotational mass moments of inertia are not equal at a joint.

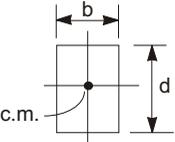
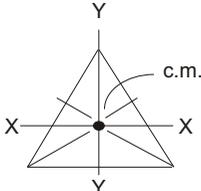
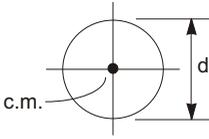
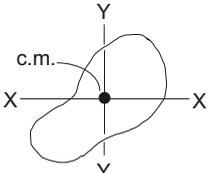
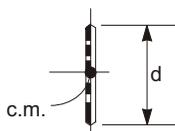
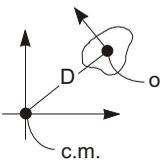
Shape in plan	Mass Moment of Inertia about vertical axis (normal to paper) through center of mass	Formula
	<p>Rectangular diaphragm: Uniformly distributed mass per unit area Total mass of diaphragm = M (or w/g)</p>	$MMI_{cm} = \frac{M(b^2+d^2)}{12}$
	<p>Triangular diaphragm: Uniformly distributed mass per unit area Total mass of diaphragm = M (or w/g)</p>	<p>Use general diaphragm formula</p>
	<p>Circular diaphragm: Uniformly distributed mass per unit area Total mass of diaphragm = M (or w/g)</p>	$MMI_{cm} = \frac{Md^2}{8}$
	<p>General diaphragm: Uniformly distributed mass per unit area Total mass of diaphragm = M (or w/g) Area of diaphragm = A Moment of inertia of area about X-X = I_x Moment of inertia of area about Y-Y = I_y</p>	$MMI_{cm} = \frac{M(I_x+I_y)}{A}$
	<p>Line mass: Uniformly distributed mass per unit length Total mass of line = M (or w/g)</p>	$MMI_{cm} = \frac{Md^2}{12}$
	<p>Axis transformation for a mass: If mass is a point mass, $MMI_O = 0$</p>	$MMI_{cm} = MMI_O + MD^2$

Figure 7
Formulae for Mass Moments of Inertia

These coupling terms will be discarded by the program, resulting in some loss of accuracy. For this reason, it is recommended that you choose joint local coordinate systems that are aligned with the principal directions of translational or rotational mass at a joint, and then specify mass values in these joint local coordinates.

Mass values must be given in consistent mass units (W/g) and mass moments of inertia must be in WL^2/g units. Here W is weight, L is length, and g is the acceleration due to gravity. The net mass values at each joint in the structure should be zero or positive.

See Figure 7 (page 41) for mass moment of inertia formulations for various planar configurations.

For more information:

- See Topic “Degrees of Freedom” (page 30) in this Chapter.
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Static and Dynamic Analysis” (page 341).

Force Load

The Force Load is used to apply concentrated forces and moments at the joints. Values may be specified in a fixed coordinate system (global or alternate coordinates) or the joint local coordinate system. All forces and moments at a joint are transformed to the joint local coordinate system and added together. The specified values are shown in Figure 8 (page 43).

Forces and moments applied along restrained degrees of freedom add to the corresponding reaction, but do not otherwise affect the structure.

For more information:

- See Topic “Degrees of Freedom” (page 30) in this Chapter.
- See Chapter “Load Patterns” (page 321).

Ground Displacement Load

The Ground Displacement Load is used to apply specified displacements (translations and rotations) at the *grounded end* of joint restraints, joint spring, and one-joint Link/Support objects. Displacements may be specified in a fixed coordinate system (global or alternate coordinates) or the joint local coordinate system.

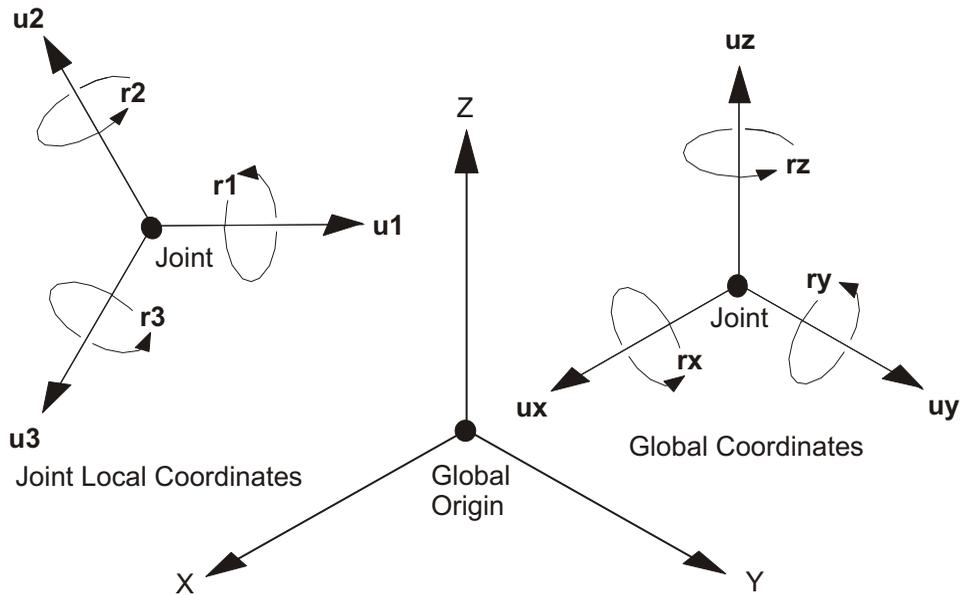


Figure 8
Specified Values for Force Load and Ground Displacement Load

The specified values are shown in Figure 8 (page 43). All displacements at a joint are transformed to the joint local coordinate system and added together.

Restraints may be considered as *rigid* connections between the joint degrees of freedom and the ground. Springs and one-joint Link/Support objects may be considered as *flexible* connections between the joint degrees of freedom and the ground.

It is very important to understand that ground displacement load applies to the ground, and does not affect the structure unless the structure is supported by restraints, springs, or one-joint Link/Supports in the direction of loading!

Restraint Displacements

If a particular joint degree of freedom is restrained, the displacement of the joint is equal to the ground displacement along that local degree of freedom. This applies regardless of whether or not springs are present.

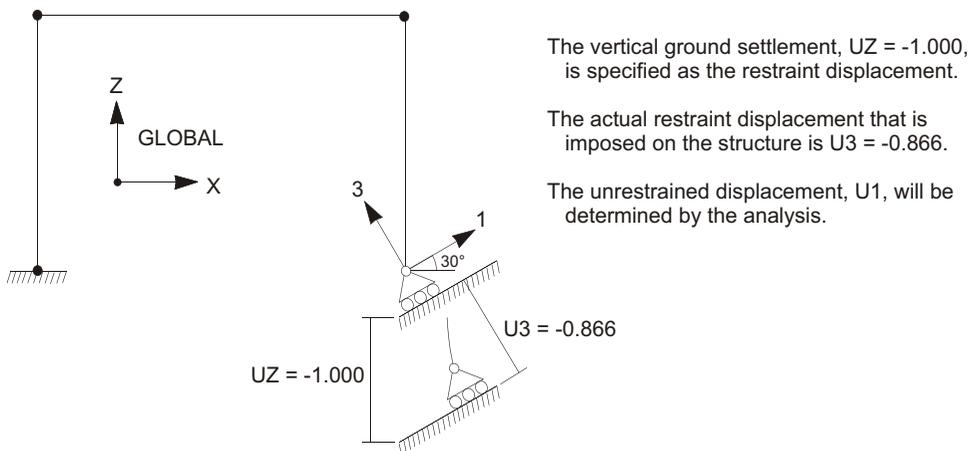


Figure 9

Example of Restraint Displacement Not Aligned with Local Degrees of Freedom

Components of ground displacement that are not along restrained degrees of freedom do not load the structure (except possibly through springs and one-joint links). An example of this is illustrated in Figure 9 (page 44).

The ground displacement, and hence the joint displacement, may vary from one Load Pattern to the next. If no ground displacement load is specified for a restrained degree of freedom, the joint displacement is zero for that Load Pattern.

Spring Displacements

The ground displacements at a joint are multiplied by the spring stiffness coefficients to obtain effective forces and moments that are applied to the joint. Spring displacements applied in a direction with no spring stiffness result in zero applied load. The ground displacement, and hence the applied forces and moments, may vary from one Load Pattern to the next.

In a joint local coordinate system, the applied forces and moments F_1 , F_2 , F_3 , M_1 , M_2 and M_3 at a joint due to ground displacements are given by:

$$\begin{array}{rcccccccc}
 F_1 & \mathbf{u1} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & u_{g1} \\
 F_2 & & \mathbf{u2} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & u_{g2} \\
 F_3 & & & \mathbf{u3} & \mathbf{0} & \mathbf{0} & \mathbf{0} & u_{g3} \\
 M_1 & & & & \mathbf{r1} & \mathbf{0} & \mathbf{0} & r_{g1} \\
 M_2 & & \mathbf{sym.} & & & \mathbf{r2} & \mathbf{0} & r_{g2} \\
 M_3 & & & & & & \mathbf{r3} & r_{g3}
 \end{array} \quad (\text{Eqn. 2})$$

where u_{g1} , u_{g2} , u_{g3} , r_{g1} , r_{g2} and r_{g3} are the ground displacements and rotations, and the terms $\mathbf{u1}$, $\mathbf{u2}$, $\mathbf{u3}$, $\mathbf{r1}$, $\mathbf{r2}$, and $\mathbf{r3}$ are the specified spring stiffness coefficients.

The net spring forces and moments acting on the joint are the sum of the forces and moments given in Equations (1) and (2); note that these are of opposite sign. At a restrained degree of freedom, the joint displacement is equal to the ground displacement, and hence the net spring force is zero.

For more information:

- See Topic “Restrains and Reactions” (page 34) in this Chapter.
- See Topic “Springs” (page 36) in this Chapter.
- See Chapter “Load Patterns” (page 321).

Link/Support Displacements

One-joint Link/Support objects are converted to zero-length, two-joint Link/Support elements. A restrained joint is generated and the ground displacement is applied as a restraint displacement at this generated joint.

The effect of the ground displacement on the structure depends upon the properties of the Link/Support element connecting the restrained joint to the structure, similar to how springs supports work, except the Link/Support stiffness may be nonlinear.

Generalized Displacements

A generalized displacement is a named displacement measure that you define. It is simply a linear combination of displacement degrees of freedom from one or more joints.

For example, you could define a generalized displacement that is the difference of the UX displacements at two joints on different stories of a building and name it

“DRIFTX”. You could define another generalized displacement that is the sum of three rotations about the Z axis, each scaled by 1/3, and name it “AVGRZ.”

Generalized displacements are primarily used for output purposes, except that you can also use a generalized displacement to monitor a nonlinear static analysis.

To define a generalized displacement, specify the following:

- A unique name
- The type of displacement measure
- A list of the joint degrees of freedom and their corresponding scale factors that will be summed to create the generalized displacement

The type of displacement measure can be one of the following:

- **Translational:** The generalized displacement scales (with change of units) as length. Coefficients of contributing joint translations are unitless. Coefficients of contributing joint rotations scale as length.
- **Rotational:** The generalized displacement is unitless (radians). Coefficients of joint translations scale as inverse length. Coefficients of joint rotations are unitless.

Be sure to choose your scale factors for each contributing component to account for the type of generalized displacement being defined.

Degree of Freedom Output

A table of the types of degrees of freedom present at every joint in the model is printed in the analysis output (.OUT) file under the heading:

DISPLACEMENT DEGREES OF FREEDOM

The degrees of freedom are listed for all of the regular joints, as well as for the master joints created automatically by the program. For Constraints, the master joints are identified by the labels of their corresponding Constraints. For Welds, the master joint for each set of joints that are welded together is identified by the label of one of the welded joints. Joints are printed in alpha-numeric order of the labels.

The type of each of the six degrees of freedom at a joint is identified by the following symbols:

- (A) Active degree of freedom
- (-) Restrained degree of freedom

- (+) Constrained degree of freedom
- () Null or unavailable degree of freedom

The degrees of freedom are always referred to the local axes of the joint. They are identified in the output as U1, U2, U3, R1, R2, and R3 for all joints. However, if all regular joints use the global coordinate system as the local system (the usual situation), then the degrees of freedom for the regular joints are identified as UX, UY, UZ, RX, RY, and RZ.

The types of degrees of freedom are a property of the structure and are independent of the Load Cases, except when staged construction is performed.

See Topic “Degrees of Freedom” (page 30) in this Chapter for more information.

Assembled Joint Mass Output

You can request assembled joint masses as part of the analysis results. The mass at a given joint includes the mass assigned directly to that joint as well as a portion of the mass from each element connected to that joint. All mass assigned to the elements is apportioned to the connected joints, so that this table represents the total mass of the structure. The masses are always referred to the local axes of the joint.

If multiple Mass Sources have been specified, the assembled joint mass output is provided for each Mass Source that was actually used in the analysis.

For more information:

- See Topic “Masses” (page 40) in this Chapter.
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Load Cases” (page 341).

Displacement Output

You can request joint displacements as part of the analysis results on a case by case basis. For dynamic Load Cases, you can also request velocities and accelerations. The output is always referred to the local axes of the joint.

- See Topic “Degrees of Freedom” (page 30) in this Chapter.
- See Chapter “Load Cases” (page 341).

Force Output

You can request joint support forces as part of the analysis results on a case by case basis. These support forces are called reactions, and are the sum of all forces from restraints, springs, or one-joint Link/Support objects at that joint. The reactions at joints that are not supported will be zero.

Note that reactions for one-joint Link/Support objects are not reported at the original joint, but rather at a generated joint at the same location with an identifying label.

The forces and moments are always referred to the local axes of the joint. The values reported are always the forces and moments that *act on* the joints. Thus a positive value of joint force or moment would tend to cause a positive value of joint translation or rotation along the corresponding degree of freedom if it were not supported.

For more information:

- See Topic “Degrees of Freedom” (page 30) in this Chapter.
- See Chapter “Load Cases” (page 341).

Element Joint Force Output

The **element joint forces** are concentrated forces and moments acting at the joints of the element that represent the effect of the rest of the structure upon the element and that cause the deformation of the element. The moments will always be zero for the solid-type elements: Plane, Asolid, and Solid.

A positive value of force or moment tends to cause a positive value of translation or rotation of the element along the corresponding joint degree of freedom.

Element joint forces must not be confused with internal forces and moments which, like stresses, act *within* the volume of the element.

For a given element, the vector of element joint forces, \mathbf{f} , is computed as:

$$\mathbf{f} = \mathbf{K} \mathbf{u} + \mathbf{r}$$

where \mathbf{K} is the element stiffness matrix, \mathbf{u} is the vector of element joint displacements, and \mathbf{r} is the vector of element applied loads as apportioned to the joints. The element joint forces are always referred to the local axes of the individual joints. They are identified in the output as F1, F2, F3, M1, M2, and M3.

Constraints and Welds

Constraints are used to enforce certain types of rigid-body behavior, to connect together different parts of the model, and to impose certain types of symmetry conditions. **Welds** are used to generate a set of constraints that connect together different parts of the model.

Basic Topics for All Users

- Overview
- Body Constraint
- Plane Definition
- Diaphragm Constraint
- Plate Constraint
- Axis Definition
- Rod Constraint
- Beam Constraint
- Equal Constraint
- Welds

Advanced Topics

- Local Constraint
- Automatic Master Joints
- Constraint Output

Overview

A constraint consists of a set of two or more **constrained joints**. The displacements of each pair of joints in the constraint are related by constraint equations. The types of behavior that can be enforced by constraints are:

- Rigid-body behavior, in which the constrained joints translate and rotate together as if connected by rigid links. The types of rigid behavior that can be modeled are:
 - Rigid Body: fully rigid for all displacements
 - Rigid Diaphragm: rigid for membrane behavior in a plane
 - Rigid Plate: rigid for plate bending in a plane
 - Rigid Rod: rigid for extension along an axis
 - Rigid Beam: rigid for beam bending on an axis
- Equal-displacement behavior, in which the translations and rotations are equal at the constrained joints
- Symmetry and anti-symmetry conditions

The use of constraints reduces the number of equations in the system to be solved and will usually result in increased computational efficiency.

Most constraint types must be defined with respect to some fixed coordinate system. The coordinate system may be the global coordinate system or an alternate coordinate system, or it may be automatically determined from the locations of the constrained joints. The Local Constraint does not use a fixed coordinate system, but references each joint using its own joint local coordinate system.

Welds are used to connect together different parts of the model that were defined separately. Each Weld consists of a set of joints that *may* be joined. The program searches for joints in each Weld that share the same location in space and constrains them to act as a single joint.

Body Constraint

A Body Constraint causes all of its constrained joints to move together as a three-dimensional rigid body. By default, all degrees of freedom at each connected joint participate. However, you can select a subset of the degrees of freedom to be constrained.

This Constraint can be used to:

- Model rigid connections, such as where several beams and/or columns frame together
- Connect together different parts of the structural model that were defined using separate meshes
- Connect Frame elements that are acting as eccentric stiffeners to Shell elements

Welds can be used to automatically generate Body Constraints for the purpose of connecting *coincident joints*.

See Topic “Welds” (page 64) in this Chapter for more information.

Joint Connectivity

Each Body Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space.

Local Coordinate System

Each Body Constraint has its own local coordinate system, the axes of which are denoted 1, 2, and 3. These correspond to the X, Y, and Z axes of a fixed coordinate system that you choose.

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts I and J) in a Body Constraint. These equations are expressed in terms of the translations (u_1 , u_2 , and u_3), the rotations (r_1 , r_2 , and r_3), and the coordinates (x_1 , x_2 , and x_3), all taken in the Constraint local coordinate system:

$$u_{1j} = u_{1i} + r_{2i} \ x_3 - r_{3i} \ x_2$$

$$u_{2j} = u_{2i} + r_{3i} \ x_1 - r_{1i} \ x_3$$

$$u_{3j} = u_{3i} + r_{1i} x_{2j} - r_{2i} x_{1j}$$

$$r_{1i} = r_{1j}$$

$$r_{2i} = r_{2j}$$

$$r_{3i} = r_{3j}$$

where $x_1 = x_{1j}$, $x_2 = x_{2j}$, and $x_3 = x_{3j}$.

If you omit any particular degree of freedom, the corresponding constraint equation is not enforced. If you omit a rotational degree of freedom, the corresponding terms are removed from the equations for the translational degrees of freedom.

Plane Definition

The constraint equations for each Diaphragm or Plate Constraint are written with respect to a particular plane. The location of the plane is not important, only its orientation.

By default, the plane is determined automatically by the program from the spatial distribution of the constrained joints as follows:

- The centroid of the constrained joints is determined
- The second moments of the locations of all of the constrained joints about the centroid are determined
- The principal values and directions of these second moments are found
- The direction of the smallest principal second moment is taken as the normal to the constraint plane; if all constrained joints lie in a unique plane, this smallest principal moment will be zero
- If no unique direction can be found, a horizontal (X-Y) plane is assumed in coordinate system **csys**; this situation can occur if the joints are coincident or collinear, or if the spatial distribution is more nearly three-dimensional than planar.

You may override automatic plane selection by specifying the following:

- **csys**: A fixed coordinate system (the default is zero, indicating the global coordinate system)
- **axis**: The axis (X, Y, or Z) normal to the plane of the constraint, taken in coordinate system **csys**.

This may be useful, for example, to specify a horizontal plane for a floor with a small step in it.

Diaphragm Constraint

A Diaphragm Constraint causes all of its constrained joints to move together as a planar diaphragm that is rigid against membrane deformation. Effectively, all constrained joints are connected to each other by links that are rigid in the plane, but do not affect out-of-plane (plate) deformation.

This Constraint can be used to:

- Model concrete floors (or concrete-filled decks) in building structures, which typically have very high in-plane stiffness
- Model diaphragms in bridge superstructures

The use of the Diaphragm Constraint for building structures eliminates the numerical-accuracy problems created when the large in-plane stiffness of a floor diaphragm is modeled with membrane elements. It is also very useful in the lateral (horizontal) dynamic analysis of buildings, as it results in a significant reduction in the size of the eigenvalue problem to be solved. See Figure 10 (page 54) for an illustration of a floor diaphragm.

Joint Connectivity

Each Diaphragm Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space, but for best results all joints should lie in the plane of the constraint. Otherwise, bending moments may be generated that are restrained by the Constraint, which unrealistically stiffens the structure. If this happens, the constraint forces reported in the analysis results may not be in equilibrium.

Local Coordinate System

Each Diaphragm Constraint has its own local coordinate system, the axes of which are denoted 1, 2, and 3. Local axis 3 is always normal to the plane of the constraint. The program arbitrarily chooses the orientation of axes 1 and 2 in the plane. The actual orientation of the planar axes is not important since only the normal direction affects the constraint equations. For more information, see Topic “Plane Definition” (page 52) in this Chapter.

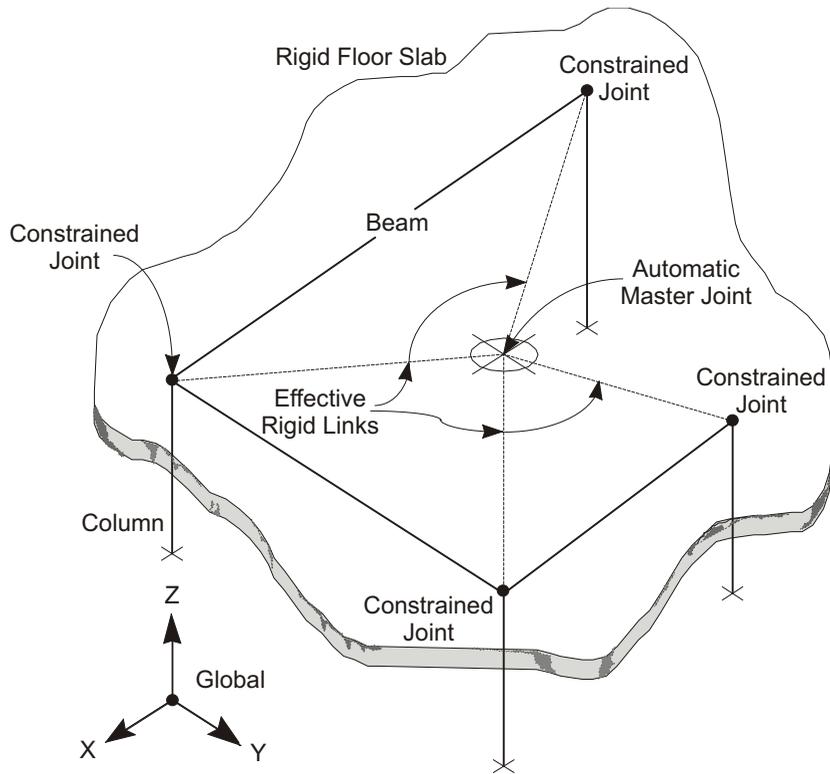


Figure 10
Use of the Diaphragm Constraint to Model a Rigid Floor Slab

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts i and j) in a Diaphragm Constraint. These equations are expressed in terms of in-plane translations (u_1 and u_2), the rotation (r_3) about the normal, and the in-plane coordinates (x_1 and x_2), all taken in the Constraint local coordinate system:

$$u_{1j} = u_{1i} - r_{3i} \ x_2$$

$$u_{2j} = u_{2i} + r_{3i} \ x_1$$

$$r_{3i} = r_{3j}$$

where $x_1 = x_{1j} - x_{1i}$ and $x_2 = x_{2j} - x_{2i}$.

Plate Constraint

A Plate Constraint causes all of its constrained joints to move together as a flat plate that is rigid against bending deformation. Effectively, all constrained joints are connected to each other by links that are rigid for out-of-plane bending, but do not affect in-plane (membrane) deformation.

This Constraint can be used to:

- Connect structural-type elements (Frame and Shell) to solid-type elements (Plane and Solid); the rotation in the structural element can be converted to a pair of equal and opposite translations in the solid element by the Constraint
- Enforce the assumption that “plane sections remain plane” in detailed models of beam bending

Joint Connectivity

Each Plate Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space. Unlike the Diaphragm Constraint, equilibrium is not affected by whether or not all joints lie in the plane of the Plate Constraint.

Local Coordinate System

Each Plate Constraint has its own local coordinate system, the axes of which are denoted 1, 2, and 3. Local axis 3 is always normal to the plane of the constraint. The program arbitrarily chooses the orientation of axes 1 and 2 in the plane. The actual orientation of the planar axes is not important since only the normal direction affects the constraint equations.

For more information, see Topic “Plane Definition” (page 52) in this Chapter.

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts I and J) in a Plate Constraint. These equations are expressed in terms of the out-of-plane translation (u_3), the bending rotations (r_1 and r_2), and the in-plane coordinates (x_1 and x_2), all taken in the Constraint local coordinate system:

$$u_{3j} = u_{3i} + r_{1i} x_2 - r_{2i} x_1$$

$$r_{1i} = r_{1j}$$

$$r_{2i} = r_{2j}$$

where $x_1 = x_{1j}$ x_{1i} and $x_2 = x_{2j}$ x_{2i} .

Axis Definition

The constraint equations for each Rod or Beam Constraint are written with respect to a particular axis. The location of the axis is not important, only its orientation.

By default, the axis is determined automatically by the program from the spatial distribution of the constrained joints as follows:

- The centroid of the constrained joints is determined
- The second moments of the locations of all of the constrained joints about the centroid are determined
- The principal values and directions of these second moments are found
- The direction of the largest principal second moment is taken as the axis of the constraint; if all constrained joints lie on a unique axis, the two smallest principal moments will be zero
- If no unique direction can be found, a vertical (Z) axis is assumed in coordinate system **csys**; this situation can occur if the joints are coincident, or if the spatial distribution is more nearly planar or three-dimensional than linear.

You may override automatic axis selection by specifying the following:

- **csys**: A fixed coordinate system (the default is zero, indicating the global coordinate system)
- **axis**: The axis (X, Y, or Z) of the constraint, taken in coordinate system **csys**.

This may be useful, for example, to specify a vertical axis for a column with a small offset in it.

Rod Constraint

A Rod Constraint causes all of its constrained joints to move together as a straight rod that is rigid against axial deformation. Effectively, all constrained joints maintain a fixed distance from each other in the direction parallel to the axis of the rod, but translations normal to the axis and all rotations are unaffected.

This Constraint can be used to:

- Prevent axial deformation in Frame elements
- Model rigid truss-like links

An example of the use of the Rod Constraint is in the analysis of the two-dimensional frame shown in Figure 11 (page 58). If the axial deformations in the beams are negligible, a single Rod Constraint could be defined containing the five joints. Instead of five equations, the program would use a single equation to define the X-displacement of the whole floor. However, it should be noted that this will result in the axial forces of the beams being output as zero, as the Constraint will cause the ends of the beams to translate together in the X-direction. Interpretations of such results associated with the use of Constraints should be clearly understood.

Joint Connectivity

Each Rod Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space, but for best results all joints should lie on the axis of the constraint. Otherwise, bending moments may be generated that are restrained by the Constraint, which unrealistically stiffens the structure. If this happens, the constraint forces reported in the analysis results may not be in equilibrium.

Local Coordinate System

Each Rod Constraint has its own local coordinate system, the axes of which are denoted 1, 2, and 3. Local axis 1 is always the axis of the constraint. The program arbitrarily chooses the orientation of the transverse axes 2 and 3. The actual orientation of the transverse axes is not important since only the axial direction affects the constraint equations.

For more information, see Topic “Axis Definition” (page 56) in this Chapter.

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts i and j) in a Rod Constraint. These equations are expressed only in terms of the axial translation (u_i):

$$u_j = u_i$$

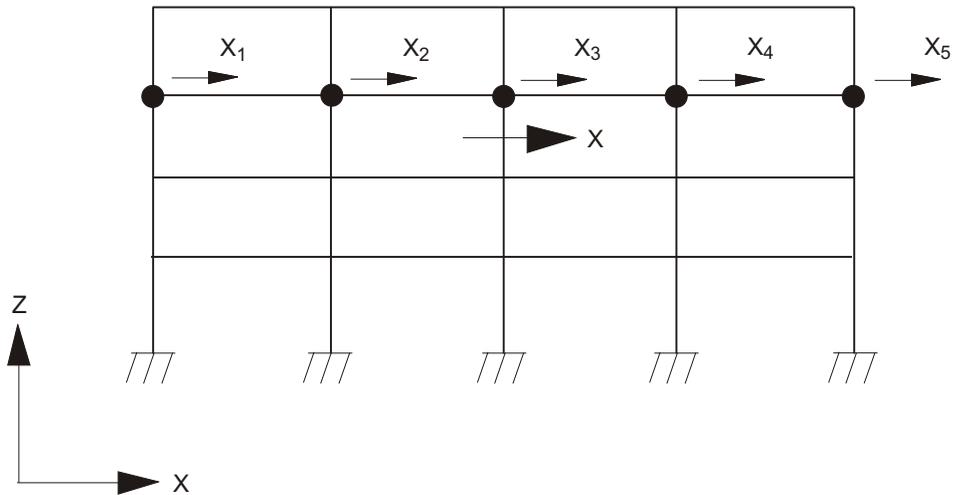


Figure 11
Use of the Rod Constraint to Model Axially Rigid Beams

Beam Constraint

A Beam Constraint causes all of its constrained joints to move together as a straight beam that is rigid against bending deformation. Effectively, all constrained joints are connected to each other by links that are rigid for off-axis bending, but do not affect translation along or rotation about the axis.

This Constraint can be used to:

- Connect structural-type elements (Frame and Shell) to solid-type elements (Plane and Solid); the rotation in the structural element can be converted to a pair of equal and opposite translations in the solid element by the Constraint
- Prevent bending deformation in Frame elements

Joint Connectivity

Each Beam Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space, but for best results all joints should lie on the axis of the constraint. Otherwise, torsional moments may be generated that are restrained by the Constraint, which unrealistically stiffens the structure. If this hap-

pens, the constraint forces reported in the analysis results may not be in equilibrium.

Local Coordinate System

Each Beam Constraint has its own local coordinate system, the axes of which are denoted 1, 2, and 3. Local axis 1 is always the axis of the constraint. The program arbitrarily chooses the orientation of the transverse axes 2 and 3. The actual orientation of the transverse axes is not important since only the axial direction affects the constraint equations.

For more information, see Topic “Axis Definition” (page 56) in this Chapter.

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts I and J) in a Beam Constraint. These equations are expressed in terms of the transverse translations (u_2 and u_3), the transverse rotations (r_2 and r_3), and the axial coordinate (x_1), all taken in the Constraint local coordinate system:

$$u_{2j} = u_{2i} + r_{3i} x_1$$

$$u_{3j} = u_{3i} - r_{2i} x_1$$

$$r_{2i} = r_{2j}$$

$$r_{3i} = r_{3j}$$

where $x_1 = x_{1j} - x_{1i}$.

Equal Constraint

An Equal Constraint causes all of its constrained joints to move together with the same displacements for each selected degree of freedom, taken in the constraint local coordinate system. The other degrees of freedom are unaffected.

The Equal Constraint differs from the rigid-body types of Constraints in that there is *no coupling* between the rotations and the translations.

This Constraint can be used to partially connect together different parts of the structural model, such as at expansion joints and hinges

For fully connecting meshes, it is better to use the Body Constraint when the constrained joints are not in exactly the same location.

Joint Connectivity

Each Equal Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space, but for best results all joints should share the same location in space if used for connecting meshes. Otherwise, moments may be generated that are restrained by the Constraint, which unrealistically stiffens the structure. If this happens, the constraint forces reported in the analysis results may not be in equilibrium.

Local Coordinate System

Each Equal Constraint uses a fixed coordinate system, **csys**, that you specify. The default for **csys** is zero, indicating the global coordinate system. The axes of the fixed coordinate system are denoted X, Y, and Z.

Selected Degrees of Freedom

For each Equal Constraint you may specify a list, **cdofs**, of up to six degrees of freedom in coordinate system **csys** that are to be constrained. The degrees of freedom are indicated as UX, UY, UZ, RX, RY, and RZ.

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts *I* and *j*) in an Equal Constraint. These equations are expressed in terms of the translations (u_x , u_y , and u_z) and the rotations (r_x , r_y , and r_z), all taken in fixed coordinate system **csys**:

$$u_{xj} = u_{xi}$$

$$u_{yj} = u_{yi}$$

$$u_{zj} = u_{zi}$$

$$r_{1i} = r_{1j}$$

$$r_{2i} = r_{2j}$$

$$r_{3i} = r_{3j}$$

If you omit any of the six degrees of freedom from the constraint definition, the corresponding constraint equation is not enforced.

Local Constraint

A Local Constraint causes all of its constrained joints to move together with the same displacements for each selected degree of freedom, taken in the separate joint local coordinate systems. The other degrees of freedom are unaffected.

The Local Constraint differs from the rigid-body types of Constraints in that there is *no coupling* between the rotations and the translations. The Local Constraint is the same as the Equal Constraint if all constrained joints have the same local coordinate system.

This Constraint can be used to:

- Model symmetry conditions with respect to a line or a point
- Model displacements constrained by mechanisms

The behavior of this Constraint is dependent upon the choice of the local coordinate systems of the constrained joints.

Joint Connectivity

Each Local Constraint connects a set of two or more joints together. The joints may have any arbitrary location in space. If the joints do not share the same location in space, moments may be generated that are restrained by the Constraint. If this happens, the constraint forces reported in the analysis results may not be in equilibrium. These moments are necessary to enforce the desired symmetry of the displacements when the applied loads are not symmetric, or may represent the constraining action of a mechanism.

For more information, see:

- Topic “Force Output” (page 48) in Chapter “Joints and Degrees of Freedom.”
- Topic “Global Force Balance Output” (page 45) in Chapter “Joints and Degrees of Freedom.”

No Local Coordinate System

A Local Constraint does not have its own local coordinate system. The constraint equations are written in terms of constrained joint local coordinate systems, which may differ. The axes of these coordinate systems are denoted 1, 2, and 3.

Selected Degrees of Freedom

For each Local Constraint you may specify a list, **ldofs**, of up to six degrees of freedom in the joint local coordinate systems that are to be constrained. The degrees of freedom are indicated as U1, U2, U3, R1, R2, and R3.

Constraint Equations

The constraint equations relate the displacements at any two constrained joints (subscripts I and J) in a Local Constraint. These equations are expressed in terms of the translations (u_1 , u_2 , and u_3) and the rotations (r_1 , r_2 , and r_3), all taken in joint local coordinate systems. The equations used depend upon the selected degrees of freedom and their signs. Some important cases are described next.

Axisymmetry

Axisymmetry is a type of symmetry about a line. It is best described in terms of a cylindrical coordinate system having its Z axis on the line of symmetry. The structure, loading, and displacements are each said to be axisymmetric about a line if they do not vary with angular position around the line, i.e., they are independent of the angular coordinate CA .

To enforce axisymmetry using the Local Constraint:

- Model any cylindrical sector of the structure using any axisymmetric mesh of joints and elements
- Assign each joint a local coordinate system such that local axes 1, 2, and 3 correspond to the coordinate directions +CR, +CA, and +CZ, respectively
- For each axisymmetric set of joints (i.e., having the same coordinates CR and CZ, but different CA), define a Local Constraint using all six degrees of freedom: U1, U2, U3, R1, R2, and R3
- Restrain joints that lie on the line of symmetry so that, at most, only axial translations (U3) and rotations (R3) are permitted

The corresponding constraint equations are:

$$u_{1j} = u_{1i}$$

$$u_{2j} = u_{2i}$$

$$u_{3j} = u_{3i}$$

$$r_{1i} = r_{1j}$$

$$r_{2i} = r_{2j}$$

$$r_{3i} = r_{3j}$$

The numeric subscripts refer to the corresponding joint local coordinate systems.

Cyclic symmetry

Cyclic symmetry is another type of symmetry about a line. It is best described in terms of a cylindrical coordinate system having its Z axis on the line of symmetry. The structure, loading, and displacements are each said to be cyclically symmetric about a line if they vary with angular position in a repeated (periodic) fashion.

To enforce cyclic symmetry using the Local Constraint:

- Model any number of adjacent, representative, cylindrical sectors of the structure; denote the size of a single sector by the angle
- Assign each joint a local coordinate system such that local axes 1, 2, and 3 correspond to the coordinate directions +CR, +CA, and +CZ, respectively
- For each cyclically symmetric set of joints (i.e., having the same coordinates CR and CZ, but with coordinate CA differing by multiples of θ), define a Local Constraint using all six degrees of freedom: U1, U2, U3, R1, R2, and R3.
- Restrain joints that lie on the line of symmetry so that, at most, only axial translations (U3) and rotations (R3) are permitted

The corresponding constraint equations are:

$$u_{1j} = u_{1i}$$

$$u_{2j} = u_{2i}$$

$$u_{3j} = u_{3i}$$

$$r_{1i} = r_{1j}$$

$$r_{2i} = r_{2j}$$

$$r_{3i} = r_{3j}$$

The numeric subscripts refer to the corresponding joint local coordinate systems.

For example, suppose a structure is composed of six identical 60° sectors, identically loaded. If two adjacent sectors were modeled, each Local Constraint would apply to a set of two joints, except that three joints would be constrained on the symmetry planes at 0°, 60°, and 120°.

If a single sector is modeled, only joints on the symmetry planes need to be constrained.

Symmetry About a Point

Symmetry about a point is best described in terms of a spherical coordinate system having its Z axis on the line of symmetry. The structure, loading, and displacements are each said to be symmetric about a point if they do not vary with angular position about the point, i.e., they are independent of the angular coordinates SB and SA. Radial translation is the only displacement component that is permissible.

To enforce symmetry about a point using the Local Constraint:

- Model any spherical sector of the structure using any symmetric mesh of joints and elements
- Assign each joint a local coordinate system such that local axes 1, 2, and 3 correspond to the coordinate directions +SB, +SA, and +SR, respectively
- For each symmetric set of joints (i.e., having the same coordinate SR, but different coordinates SB and SA), define a Local Constraint using only degree of freedom U3
- For all joints, restrain the degrees of freedom U1, U2, R1, R2, and R3
- Fully restrain any joints that lie at the point of symmetry

The corresponding constraint equations are:

$$u_{3j} = u_{3i}$$

The numeric subscripts refer to the corresponding joint local coordinate systems.

It is also possible to define a case for symmetry about a point that is similar to cyclic symmetry around a line, e.g., where each octant of the structure is identical.

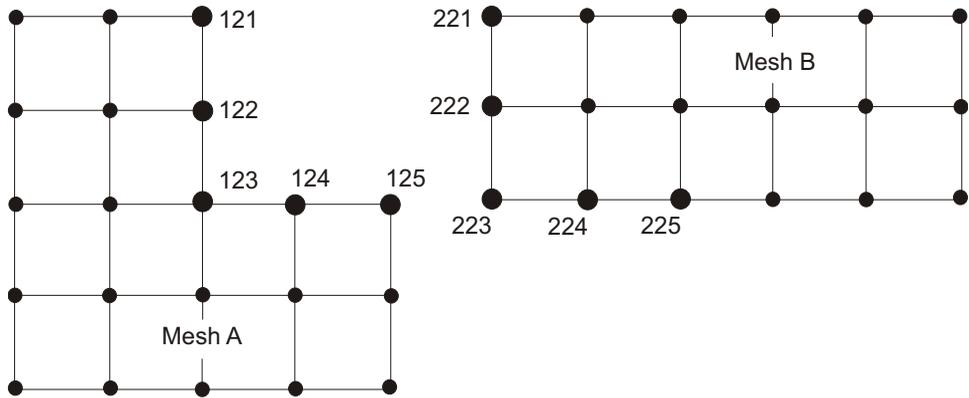


Figure 12
Use of a Weld to Connect Separate Meshes at Coincident Joints

Welds

A Weld can be used to connect together different parts of the structural model that were defined using separate meshes. A Weld is not a single Constraint, but rather is a set of joints from which the program will automatically generate multiple Body Constraints to connect together coincident joints.

Joints are considered to be coincident if the distance between them is less than or equal to a tolerance, **tol**, that you specify. Setting the tolerance to zero is permissible but is not recommended.

One or more Welds may be defined, each with its own tolerance. Only the joints within each Weld will be checked for coincidence with each other. In the most common case, a single Weld is defined that contains all joints in the model; all coincident groups of joints will be welded. However, in situations where structural discontinuity is desired, it may be necessary to prevent the welding of some coincident joints. This may be facilitated by the use of multiple Welds.

Figure 12 (page 65) shows a model developed as two separate meshes, A and B. Joints 121 through 125 are associated with mesh A, and Joints 221 through 225 are associated with mesh B. Joints 121 through 125 share the same location in space as Joints 221 through 225, respectively. These are the interfacing joints between the two meshes. To connect these two meshes, a single Weld can be defined containing all joints, or just joints 121 through 125 and 221 through 225. The program would

generate five Body Constraints, each containing two joints, resulting in an integrated model.

It is permissible to include the same joint in more than one Weld. This could result in the joints in different Welds being constrained together if they are coincident with the common joint. For example, suppose that Weld 1 contained joints 1, 2, and 3, Weld 2 contained joints 3, 4, and 5. If joints 1, 3, and 5 were coincident, joints 1 and 3 would be constrained by Weld 1, and joints 3 and 5 would be constrained by Weld 2. The program would create a single Body Constraint containing joints 1, 3, and 5. On the other hand, if Weld 2 did not contain joint 3, the program would only generate a Body Constraint containing joint 1 and 3 from Weld 1; joint 5 would not be constrained.

For more information, see Topic “Body Constraint” (page 51) in this Chapter.

Automatic Master Joints

The program automatically creates an internal **master joint** for each explicit Constraint, and a master joint for each internal Body Constraint that is generated by a Weld. Each master joint governs the behavior of the corresponding constrained joints. The displacement at a constrained degree of freedom is computed as a linear combination of the displacements of the master joint.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom” for more information.

Stiffness, Mass, and Loads

Joint local coordinate systems, springs, masses, and loads may all be applied to constrained joints. Elements may also be connected to constrained joints. The joint and element stiffnesses, masses and loads from the constrained degrees of freedom are automatically transferred to the master joint in a consistent fashion.

The translational stiffness at the master joint is the sum of the translational stiffnesses at the constrained joints. The same is true for translational masses and loads.

The rotational stiffness at a master joint is the sum of the rotational stiffnesses at the constrained degrees of freedom, plus the second moment of the translational stiffnesses at the constrained joints for the Body, Diaphragm, Plate, and Beam Constraints. The same is true for rotational masses and loads, except that only the first moment of the translational loads is used. The moments of the translational

stiffnesses, masses, and loads are taken about the center of mass of the constrained joints. If the joints have no mass, the centroid is used.

Local Coordinate Systems

Each master joint has *two* local coordinate systems: one for the translational degrees of freedom, and one for the rotational degrees of freedom. The axes of each local system are denoted 1, 2, and 3. For the Local Constraint, these axes correspond to the local axes of the constrained joints. For other types of Constraints, these axes are chosen to be the principal directions of the translational and rotational masses of the master joint. Using the principal directions eliminates coupling between the mass components in the master-joint local coordinate system.

For a Diaphragm or Plate Constraint, the local 3 axes of the master joint are always normal to the plane of the Constraint. For a Beam or Rod Constraint, the local 1 axes of the master joint are always parallel to the axis of the Constraint.

Constraint Output

For each Body, Diaphragm, Plate, Rod, and Beam Constraint having more than two constrained joints, the following information about the Constraint and its master joint is printed in the output file:

- The translational and rotational local coordinate systems for the master joint
- The total mass and mass moments of inertia for the Constraint that have been applied to the master joint
- The center of mass for each of the three translational masses

The degrees of freedom are indicated as U1, U2, U3, R1, R2, and R3. These are referred to the two local coordinate systems of the master joint.

Material Properties

The Materials are used to define the mechanical, thermal, and density properties used by the Frame, Cable, Tendon, Shell, Plane, Asolid, and Solid elements.

Basic Topics for All Users

- Overview
- Local Coordinate System
- Stresses and Strains
- Isotropic Materials
- Uniaxial Materials
- Mass Density
- Weight Density
- Design-Type Indicator

Advanced Topics

- Orthotropic Materials
- Anisotropic Materials
- Temperature-Dependent Materials

- Element Material Temperature
- Material Damping
- Nonlinear Material Behavior
- Hysteresis Models
- Modified Darwin-Pecknold Concrete Model
- Time-dependent Properties

Overview

The Material properties may be defined as isotropic, orthotropic or anisotropic. How the properties are actually utilized depends on the element type. Each Material that you define may be used by more than one element or element type. For each element type, the Materials are referenced indirectly through the Section properties appropriate for that element type.

All elastic material properties may be temperature dependent. Properties are given at a series of specified temperatures. Properties at other temperatures are obtained by linear interpolation.

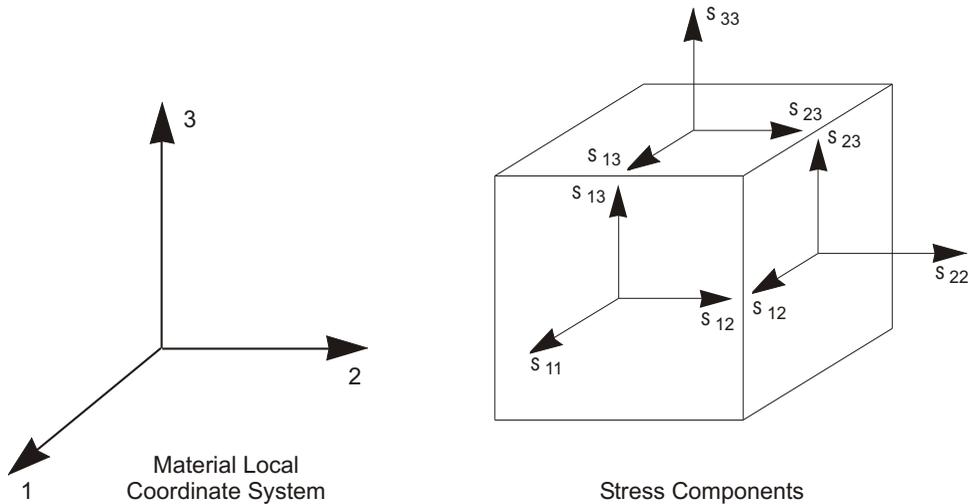
For a given execution of the program, the properties used by an element are assumed to be constant regardless of any temperature changes experienced by the structure. Each element may be assigned a **material temperature** that determines the material properties used for the analysis.

Time-dependent properties include creep, shrinkage, and age-dependent elasticity. These properties can be activated during a staged-construction analysis, and form the basis for subsequent analyses.

Nonlinear stress-strain curves may be defined for use with fiber hinges in frame elements or nonlinear layers in shell elements.

Local Coordinate System

Each Material has its own **Material local coordinate system** used to define the elastic and thermal properties. This system is significant only for orthotropic and anisotropic materials. Isotropic materials are independent of any particular coordinate system.

**Figure 13**

Definition of Stress Components in the Material Local Coordinate System

The axes of the Material local coordinate system are denoted 1, 2, and 3. By default, the Material coordinate system is aligned with the local coordinate system for each element. However, you may specify a set of one or more **material angles** that rotate the Material coordinate system with respect to the element system for those elements that permit orthotropic or anisotropic properties.

For more information:

- See Topic “Material Angle” (page 193) in Chapter “The Shell Element.”
- See Topic “Material Angle” (page 219) in Chapter “The Plane Element.”
- See Topic “Material Angle” (page 229) in Chapter “The Asolid Element.”
- See Topic “Material Angles” (page 246) in Chapter “The Solid Element.”

Stresses and Strains

The elastic mechanical properties relate the behavior of the stresses and strains within the Material. The stresses are defined as forces per unit area acting on an elemental cube aligned with the material axes as shown in Figure 13 (page 71). The stresses s_{11} , s_{22} , and s_{33} are called the direct stresses and tend to cause length

change, while τ_{12} , τ_{13} , and τ_{23} are called the shear stresses and tend to cause angle change.

Not all stress components exist in every element type. For example, the stresses τ_{22} , τ_{33} , and τ_{23} are assumed to be zero in the Frame element, and stress τ_{33} is taken to be zero in the Shell element.

The direct strains ϵ_{11} , ϵ_{22} , and ϵ_{33} measure the change in length along the Material local 1, 2, and 3 axes, respectively, and are defined as:

$$\epsilon_{11} = \frac{du_1}{dx_1}$$

$$\epsilon_{22} = \frac{du_2}{dx_2}$$

$$\epsilon_{33} = \frac{du_3}{dx_3}$$

where u_1 , u_2 , and u_3 are the displacements and x_1 , x_2 , and x_3 are the coordinates in the Material 1, 2, and 3 directions, respectively.

The engineering shear strains γ_{12} , γ_{13} , and γ_{23} , measure the change in angle in the Material local 1-2, 1-3, and 2-3 planes, respectively, and are defined as:

$$\gamma_{12} = \frac{du_1}{dx_2} + \frac{du_2}{dx_1}$$

$$\gamma_{13} = \frac{du_1}{dx_3} + \frac{du_3}{dx_1}$$

$$\gamma_{23} = \frac{du_2}{dx_3} + \frac{du_3}{dx_2}$$

Note that the engineering shear strains are equal to twice the tensorial shear strains ϵ_{12} , ϵ_{13} , and ϵ_{23} , respectively.

Strains can also be caused by a temperature change, T , that can be specified as a load on an element. No stresses are caused by a temperature change unless the induced thermal strains are restrained.

See Cook, Malkus, and Plesha (1989), or any textbook on elementary mechanics.

Orthotropic Materials

The behavior of an orthotropic material can be different in each of the three local coordinate directions. However, like an isotropic material, shearing behavior is uncoupled from extensional behavior and is not affected by temperature change.

The orthotropic mechanical and thermal properties relate strain to stress and temperature change as follows:

$$\begin{array}{ccccccc}
 & \frac{1}{\mathbf{e1}} & \frac{-\mathbf{u12}}{\mathbf{e2}} & \frac{-\mathbf{u13}}{\mathbf{e3}} & 0 & 0 & 0 & & & \\
 & & \frac{1}{\mathbf{e2}} & \frac{-\mathbf{u23}}{\mathbf{e3}} & & & & & & \\
 11 & & & & 0 & 0 & 0 & 11 & \mathbf{a1} & \\
 22 & & & \frac{1}{\mathbf{e3}} & 0 & 0 & 0 & 22 & \mathbf{a2} & \\
 33 & & & & & & & 33 & \mathbf{a3} & \\
 12 & & & & \frac{1}{\mathbf{g12}} & 0 & 0 & 12 & 0 & T \\
 13 & & & & & \frac{1}{\mathbf{g13}} & 0 & 13 & 0 & \\
 23 & \text{sym.} & & & & & & 23 & 0 & \\
 & & & & & & \frac{1}{\mathbf{g23}} & & &
 \end{array} \tag{Eqn. 3}$$

where $\mathbf{e1}$, $\mathbf{e2}$, and $\mathbf{e3}$ are the moduli of elasticity; $\mathbf{u12}$, $\mathbf{u13}$, and $\mathbf{u23}$ are the Poisson's ratios; $\mathbf{g12}$, $\mathbf{g13}$, and $\mathbf{g23}$ are the shear moduli; and $\mathbf{a1}$, $\mathbf{a2}$, and $\mathbf{a3}$ are the coefficients of thermal expansion.

Note that the elastic moduli and the shear moduli must be positive. The Poisson's ratios may take on any values provided that the upper-left 3x3 portion of the stress-strain matrix is positive-definite (i.e., has a positive determinant.)

Anisotropic Materials

The behavior of an anisotropic material can be different in each of the three local coordinate directions. In addition, shearing behavior can be fully coupled with extensional behavior and can be affected by temperature change.

The anisotropic mechanical and thermal properties relate strain to stress and temperature change as follows:

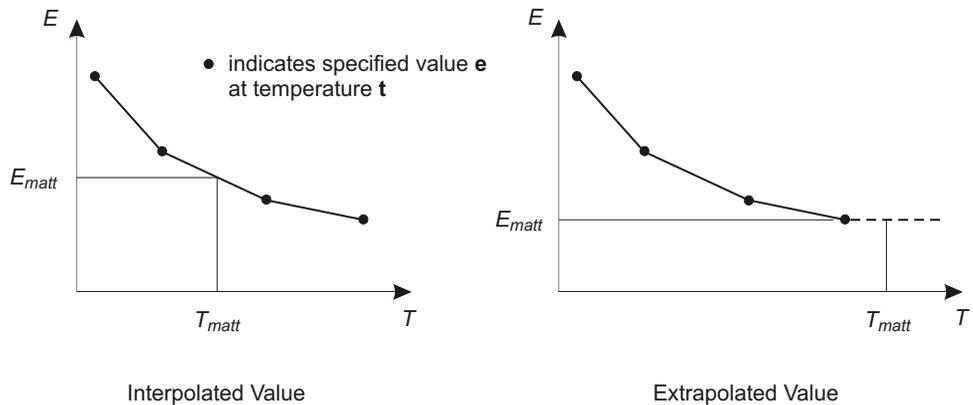


Figure 14
Determination of Property E_{matt} at Temperature T_{matt} from Function $E(T)$

Element Material Temperature

You can assign each element an **element material temperature**. This is the temperature at which temperature-dependent material properties are evaluated for the element. The properties at this fixed temperature are used for all analyses regardless of any temperature changes experienced by the element during loading.

The element material temperature may be uniform over an element or interpolated from values given at the joints. In the latter case, a uniform material temperature is used that is the average of the joint values. The default material temperature for any element is zero.

The properties for a temperature-independent material are constant regardless of the element material temperatures specified.

Mass Density

For each Material you may specify a mass density, **m**, that is used for calculating the mass of the element. The total mass of the element is the product of the mass density (mass per unit volume) and the volume of the element. This mass is apportioned to each joint of the element. The same mass is applied along of the three translational degrees of freedom. No rotational mass moments of inertia are computed.

Consistent mass units must be used. Typically the mass density is equal to the weight density of the material divided by the acceleration due to gravity, but this is not required.

The mass density property is independent of temperature.

For more information:

- See Topic “Mass” (page 134) in Chapter “The Frame Element.”
- See Topic “Mass” (page 172) in Chapter “The Cable Element.”
- See Topic “Mass” (page 206) in Chapter “The Shell Element.”
- See Topic “Mass” (page 220) in Chapter “The Plane Element.”
- See Topic “Mass” (page 232) in Chapter “The Asolid Element.”
- See Topic “Mass” (page 248) in Chapter “The Solid Element.”
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.

Weight Density

For each Material you may specify a weight density, w , that is used for calculating the self-weight of the element. The total weight of the element is the product of the weight density (weight per unit volume) and the volume of the element. This weight is apportioned to each joint of the element. Self-weight is activated using Self-weight Load and Gravity Load.

The weight density property is independent of temperature.

For more information:

- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Material Damping

You may specify material damping to be used in dynamic analyses. Different types of damping are available for different types of Load Cases. Material damping is a property of the material and affects all Load Cases of a given type in the same way. You may specify additional damping in each Load Case.

Because damping has such a significant affect upon dynamic response, you should use care in defining your damping parameters.

The material used for calculating damping is determined for the different types of elements as follows:

- For frame elements, if a material overwrite is specified, that material is used. Otherwise the material for the current frame section is used, with non-prismatic sections using a simple average of the damping coefficients over all materials along the full length of the section.
- For shell elements, if a material overwrite is specified, that material is used. Otherwise the material for the current shell section is used, with layered sections using a thickness-weighted average of the damping coefficients over all material layers in the section.
- For cable, tendon, plane, asolid, and solid elements, the material for the current section is used.

Material-based damping does not apply to link elements.

Modal Damping

The material modal damping available in SAP2000 is stiffness weighted, and is also known as **composite modal damping**. It is used for all response-spectrum and modal time-history analyses. For each material you may specify a material modal damping ratio, r , where $0 < r < 1$. The damping ratio, r_{ij} , contributed to mode I by element j of this material is given by

$$r_{ij} = \frac{r_i \mathbf{K}_{j \ i}}{K_i}$$

where $\mathbf{K}_{j \ i}$ is mode shape for mode I , \mathbf{K}_j is the stiffness matrix for element j , and K_i is the modal stiffness for mode I given by

$$K_i = \sum_j \mathbf{K}_{j \ i}^T \mathbf{K}_j \mathbf{K}_{j \ i}$$

summed over all elements, j , in the model.

Viscous Proportional Damping

Viscous proportional damping is used for direct-integration time-history analyses. For each material, you may specify a mass coefficient, c_M , and a stiffness coefficient, c_K . You may specify these two coefficients directly, or they may be computed by specifying equivalent fractions of critical modal damping at two different periods or frequencies.

The damping matrix for element j of the material is computed as:

$$\mathbf{C}_j = c_M \mathbf{M}_j + c_K \mathbf{K}_j^0$$

where \mathbf{M}_j is the mass of the element, and \mathbf{K}_j^0 is the stiffness of the element. The superscript “0” indicates that for nonlinear elements, the initial stiffness is used. This is the stiffness of the element at zero initial conditions, regardless of the current nonlinear state of the element. The exception to this rule is that if the current nonlinear state has zero stiffness *and* zero force or stress (such as for cracked concrete material), then zero damping is assumed. In the case where the initial stiffness is different in the negative and positive direction of loading, the larger stiffness is used.

Hysteretic Proportional Damping

Hysteretic proportional damping is used for steady-state and power-spectral-density analyses. For each material, you may specify a mass coefficient, d_M , and a stiffness coefficient, d_K . The hysteretic damping matrix for element j of the material is computed as:

$$\mathbf{D}_j = d_M \mathbf{M}_j + d_K \mathbf{K}_j^0$$

where \mathbf{M}_j is the mass of the element, and \mathbf{K}_j^0 is the stiffness of the element. See the subtopic “Viscous Proportional Damping” above for how the material stiffness is determined for nonlinear elements.

Nonlinear Material Behavior

Nonlinear material behavior is available in certain elements using a directional material model, in which uncoupled stress-strain behavior is modeled for one or more stress-strain components. This is a simple and practical engineering model suitable for many applications such as beams and columns, shear walls, bridge decks, tunnels, retaining walls, and others. You should carefully examine the applicability of

this model before using it in a general continuum model where the governing stresses change direction substantially from place to place.

In addition, a two-dimensional concrete model is available for use in the layered shell. This model is discussed in topic “Modified Darwin-Pecknold Concrete Model” later in this chapter. The remainder of this present topic concerns the directional material model.

Nonlinear material behavior is currently not temperature-dependent. The behavior specified at the initial (most negative) temperature is used for all material temperatures.

Tension and Compression

For each material you may specify an axial stress-strain curve that is used to represent the direct (tension-compression) stress-strain behavior of the material along any material axis. For Uniaxial materials, this represents the relationship between σ_{11} and ϵ_{11} . For Isotropic, Orthotropic, and Anisotropic materials, this curve represents the behavior along each of the three material axes, σ_{11} ϵ_{11} , σ_{22} ϵ_{22} , and σ_{33} ϵ_{33} . The nonlinear stress-strain behavior is the same in each direction, even for Orthotropic, and Anisotropic materials.

Tension is always positive, regardless of the type of material (steel, concrete, etc.) The tensile and compressive sides of the stress-strain behavior may be different from each other. For what follows, the direct stress-strain curve may be written as follows:

$$\epsilon_{ii}(\sigma_{ii}) = \begin{cases} T(\sigma_{ii}), & \sigma_{ii} \geq 0 \\ C(\sigma_{ii}), & \sigma_{ii} < 0 \end{cases} \quad (\text{Eqn. 5a})$$

where $T(\sigma)$ represents tensile behavior, and $C(\sigma)$ represents compressive behavior, subject to the restrictions:

$$\begin{aligned} T(\sigma) &\geq 0, & T(0) &= 0 \\ C(\sigma) &\leq 0, & C(0) &= 0 \end{aligned} \quad (\text{Eqn. 5b})$$

Shear

A shear stress-strain curve is computed internally from the direct stress-strain curve. The assumption is made that shearing behavior can be computed from tensile and compressive behavior acting at 45° to the material axes using Mohr's circle

in the plane. For an Isotropic, Orthotropic, or Anisotropic material, this results in the following symmetrical relationship for shear:

$$\epsilon_{ij}(\sigma_{ij}) = \begin{pmatrix} s(\sigma_{ij}), & ij & 0 \\ s(\sigma_{ij}), & ij & 0 \end{pmatrix} \quad (\text{Eqn. 6a})$$

where

$$s(\sigma_{ij}) = \frac{1}{4}(T(\sigma_{ij}) + C(\sigma_{ij})), \quad ij = \frac{1}{2} \quad ij = 0, \quad i = j \quad (\text{Eqn. 6b})$$

For the case where the direct stress-strain curve is symmetrical, such as for steel, we have $T(\sigma) = C(\sigma)$, and therefore:

$$\epsilon_{ij}(\sigma_{ij}) = \frac{1}{2} \begin{pmatrix} T(\sigma_{ij}), & ij & 0 \\ C(\sigma_{ij}), & ij & 0 \end{pmatrix}$$

To create a material where the shearing relationship, $s(\sigma)$, is primary and known, you can define a symmetrical direct stress-strain relationship such that:

$$T(\sigma) = C(\sigma) = 2s(\sigma)$$

When shear stress is considered for a Uniaxial material, the stress is half that for an Isotropic material. In this way, if you have two uniaxial materials at 90° to each other, the shear behavior matches that of an Isotropic material. Thus for Uniaxial materials:

$$\epsilon_{ij}(\sigma_{ij}) = \frac{1}{2} \begin{pmatrix} s(\sigma_{ij}), & ij & 0 \\ s(\sigma_{ij}), & ij & 0 \end{pmatrix} \quad (\text{for Uniaxial material}) \quad (\text{Eqn. 6c})$$

Hysteresis

Several hysteresis models are available to define the nonlinear stress-strain behavior when load is reversed or cycled. For the most part, these models differ in the amount of energy they dissipate in a given cycle of deformation, and how the energy dissipation behavior changes with an increasing amount of deformation.

Details are provided in Topic "Hysteresis Models" (page 85) in this chapter.

Application

Nonlinear stress-strain curves are currently used in the two applications described in the following.

Fiber Hinges

Fiber hinges are used to define the coupled axial force and bi-axial bending behavior at locations along the length of a frame element. The hinges can be defined manually, or created automatically for certain types of frame sections, including Section-Designer sections.

For each fiber in the cross section at a fiber hinge, the material direct nonlinear stress-strain curve is used to define the axial σ_{11} - ϵ_{11} relationship. Summing up the behavior of all the fibers at a cross section and multiplying by the hinge length gives the axial force-deformation and biaxial moment-rotation relationships.

The σ_{11} - ϵ_{11} is the same whether the material is Uniaxial, Isotropic, Orthotropic, or Anisotropic. Shear behavior is not considered in the fibers. Instead, shear behavior is computed for the frame section as usual using the linear shear modulus **g12**.

For more information:

- See Topic “Section Designer Sections” (page 134) in Chapter “The Frame Element.”
- See Chapter “Frame Hinge Properties” (page 147).

Layered Shell Element

The Shell element with the layered section property may consider linear, nonlinear, or mixed material behavior. For each layer, you select a material, a material angle, and whether each of the in-plane stress-strain relationships are linear, nonlinear, or inactive (zero stress). These relationships include σ_{11} - ϵ_{11} , σ_{22} - ϵ_{22} , and σ_{12} - ϵ_{12} .

For Uniaxial materials, the stress $\sigma_{22} = 0$ always. However, shear stiffness is assumed to be present, but may be set to zero by setting the shear relationship to be inactive.

For all materials, the transverse normal stress $\sigma_{33} = 0$. The transverse shear behavior is always linear, using the appropriate shear moduli **g13** and **g23** from the material matrix (Eqns. 1–4 above.)

If all three in-plane relationships for a given layer are linear, the corresponding linear matrix is used (Eqns. 1–4 above), adjusted for the plane-stress condition $\nu_{33} = 0$. Poisson effects are included, which may couple the two direct stresses.

If any of the in-plane relationships for a given layer are nonlinear or inactive, then all three relationships become uncoupled according to these rules:

- Poisson’s ratio is taken to be zero.
- Linear direct stress-strain relationships use stiffness **e1** from the material matrix (Eqns. 1-4).
- Linear shear stress-strain relationships use shear modulus **g12** (Eqns. 1-4).
- Nonlinear direct stress-strain relationships use Eqns. 5 above.
- Nonlinear shear stress-strain relationships use Eqns. 6 above.
- Inactive stress-strain relationships assume that the corresponding stress is zero.

The stress-strain behavior for a given layer is always defined in the material coordinate system specified by the material angle for that layer. It is particularly important to keep this in mind when using Uniaxial materials, for which $\theta = 0$.

The above description is for the directional material. In addition, a two-dimensional concrete model is available for use in the layered shell. This model is discussed in topic “Modified Darwin-Pecknold Concrete Model” later in this chapter.

For more information, see Subtopic “Layered Section Property” (page 193) in Chapter “The Shell Element.”

Friction and Dilatational Angles

For concrete materials, you can specify a friction angle and a dilatational angle. *These should normally be set to zero. The friction angle is an experimental parameter, and is not recommended for normal use. The dilatational angle is a future parameter, and has no effect on the model.*

The friction angle, ϕ , takes values $0 \leq \phi \leq 90$. For the recommended value of $\phi = 0$, shear behavior is as described above. For experimental use with $\phi > 0$, the shear stress is computed primarily using a frictional model having linear stiffness **g12** up to a limiting stress given by:

$$\tau_{12} = \begin{cases} \sigma \tan \phi, & \sigma \leq \sigma_c \\ \sigma_c, & \sigma > \sigma_c \end{cases} \quad (\text{Eqn. 7a})$$

where $\frac{1}{2} \epsilon_{11} - \frac{1}{2} \epsilon_{22}$). Compression is required in order to develop any shear strength by this equation. In addition, cohesion is added using (Eqn. 6a) above, but considering only the contribution to shear due to tension:

$$s(\epsilon_{12}) = \frac{1}{4} T(\epsilon_{12}), \quad \epsilon_{12} \geq \frac{1}{2} \epsilon_{12} = 0 \quad (\text{Eqn. 7b})$$

This behavior, while interesting, can present computational challenges unless the model is well defined and reasonably loaded, with sufficient ductility provided using reinforcing steel. *To repeat: This is an experimental behavior and is not recommended for normal use.*

Hysteresis Models

Hysteresis is the process of energy dissipation through deformation (displacement), as opposed to viscosity which is energy dissipation through deformation rate (velocity). Hysteresis is typical of solids, whereas viscosity is typical of fluids, although this distinction is not rigid.

Hysteretic behavior may affect nonlinear static and nonlinear time-history load cases that exhibit load reversals and cyclic loading. Monotonic loading is not affected.

Several different hysteresis models are available to describe the behavior of different types of materials. For the most part, these differ in the amount of energy they dissipate in a given cycle of deformation, and how the energy dissipation behavior changes with an increasing amount of deformation.

Each hysteresis model may be used for the following purposes:

- Material stress-strain behavior, affecting frame fiber hinges and layered shells that use the material
- Single degree-of-freedom frame hinges, such as M3 or P hinges. Interacting hinges, such as P-M3 or P-M2-M3, currently use the isotropic model
- Link/support elements of type multi-linear plasticity.

Although the present chapter concerns material properties, this discussion pertains equally to all three of these applications.

Backbone Curve (Action vs. Deformation)

For each material, hinge, or link degree of freedom, a uniaxial action vs. deformation curve defines the nonlinear behavior under monotonic loading in the positive and negative directions.

Here **action** and **deformation** are an energy conjugate pair as follows:

- For materials, stress vs. strain
- For hinges and multi-linear links, force vs. deformation or moment vs. rotation, depending upon the degree of freedom to which it is applied

For each model, the uniaxial action-deformation curve is given by a set of points that you define. This curve is called the **backbone curve**, and it can take on almost any shape, with the following restrictions:

- One point must be the origin, (0,0)
- At least one point with positive deformation, and one point with negative deformation, must be defined
- The deformations of the specified points must increase monotonically, with no two values being equal
- The action at each point must have the same sign as the deformation (they can be zero)
- The slope given by the last two points specified on the positive deformation axis is extrapolated to infinite positive deformation, or until it reaches zero value. Similarly, the slope given by the last two points specified on the negative deformation axis is extrapolated to infinite negative deformation, or until it reaches zero value.

The given curve defines the action-deformation relationship under monotonic loading. The first slope on either side of the origin is elastic; the remaining segments define plastic deformation. If the deformation reverses, it typically follows the two elastic segments before beginning plastic deformation in the reverse direction, except as described below.

Cyclic Behavior

Several hysteresis models are available in SAP2000, ETABS, and CSiBridge. The available models may vary from product to product, and may include any or all of the models described below.

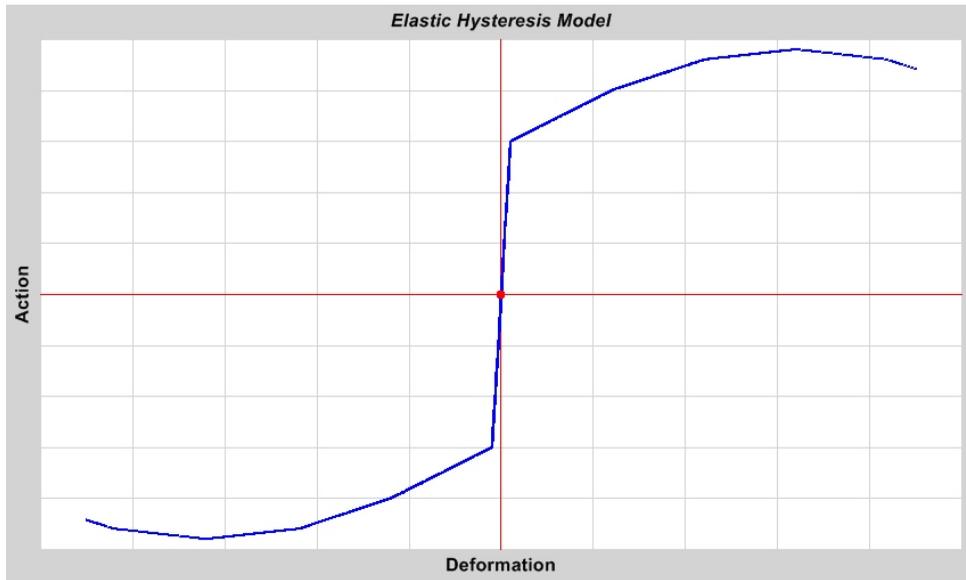


Figure 15

*Elastic Hysteresis Model under Increasing Cyclic Load No Energy Dissipation
Showing the Backbone Curve Used for All Hysteresis Figures*

Typical for all models, cyclic loading behaves as follows:

- Initial loading in the positive or negative direction follows the backbone curve
- Upon reversal of deformation, unloading occurs along a different path, usually steeper than the loading path. This is often parallel or nearly parallel to the initial elastic slope.
- After the load level is reduced to zero, continued reversal of deformation causes reverse loading along a path that eventually joins the backbone curve on the opposite side, usually at a deformation equal to the maximum previous deformation in that direction or the opposite direction.

In the descriptions below of cyclic deformation, “loading” refers to increasing the magnitude of deformation in a given positive or negative direction, and “unloading” refers to subsequent reduction of the deformation until the force level reaches zero. Continued reduction of the deformation is “reverse loading” until the deformation reaches zero, after which the deformation increases again with the same sign as the load and is “loading” again. Loading and unloading occur in the positive (first and third) quadrants of the action-deformation plot, and reverse loading occurs in the negative (second and fourth) quadrants.

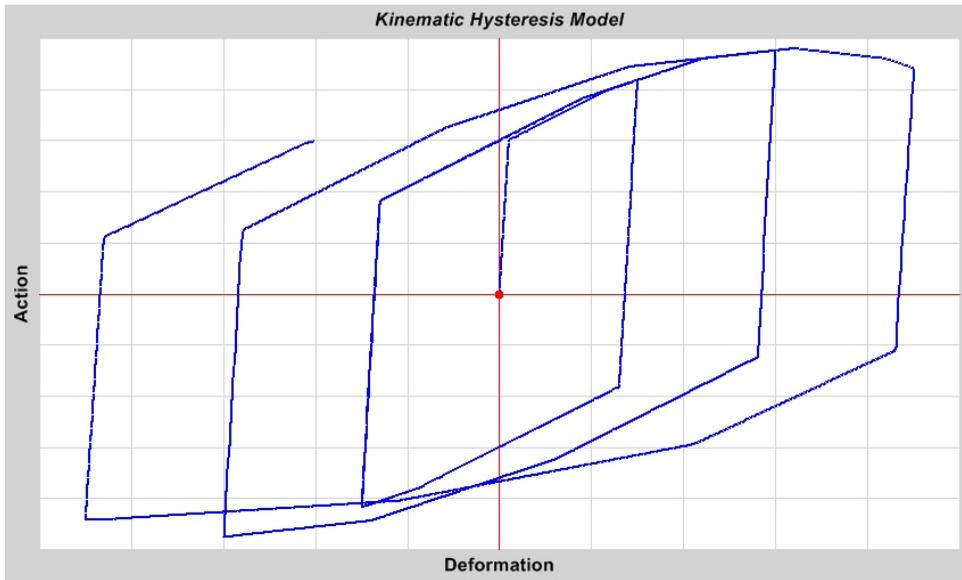


Figure 16
Kinematic Hysteresis Model under Increasing Cyclic Load

Elastic Hysteresis Model

The behavior is nonlinear but it is elastic. This means that the material always loads and unloads along the backbone curve, and no energy is dissipated. This behavior is illustrated in Figure 15 (page 87). This same backbone curve is used in the figures for all subsequent models, except that the concrete model uses only the positive portion of the curve, with the negative portion being defined separately.

Kinematic Hysteresis Model

This model is based upon kinematic hardening behavior that is commonly observed in metals, and it is the default hysteresis model for all metal materials in the program. This model dissipates a significant amount of energy, and is appropriate for ductile materials.

Under the rules of kinematic hardening, plastic deformation in one direction “pulls” the curve for the other direction along with it. Matching pairs of points are linked. No additional parameters are required for this model.

Upon unloading and reverse loading, the curve follows a path made of segments parallel to and of the same length as the previously loaded segments and their opposite-direction counterparts until it rejoins the backbone curve when loading in the opposite direction. This behavior is shown in Figure 16 (page 88) for cycles of increasing deformation.

When you define the points on the multi-linear curve, you should be aware that symmetrical pairs of points will be linked, even if the curve is not symmetrical. This gives you some control over the shape of the hysteretic loop.

The kinematic model forms the basis for several of the other model described below, including Takeda, degrading, and BRB hardening.

Degrading Hysteresis Model

This model is very similar to the Kinematic model, but uses a degrading hysteretic loop that accounts for decreasing energy dissipation and unloading stiffness with increasing plastic deformation.

Two measures are used for plastic deformation:

- Maximum plastic deformation in each the positive and negative directions
- Accumulated plastic deformation, which is the absolute sum of each increment of positive or negative plastic deformation. Plastic deformation is that which does not occur on the two elastic segments of the action-deformation curve

Accumulated plastic deformation can occur under cyclic loading of constant amplitude, and can be used to represent fatigue.

For this model, the following parameters are required:

- Separately for positive and negative deformations
 - Initial energy factor at yield, f_0 , usually 1.0
 - Energy factor at moderate deformation, f_1
 - Energy factor at maximum deformation, f_2
 - Moderate deformation level, x_1 , as a multiple of the yield deformation
 - Maximum deformation level, x_2 , as a multiple of the yield deformation
 - Accumulated deformation weighting factor, a
- Stiffness degradation weighting factor, s
- Larger-smaller weighting factor, w , usually 0.0

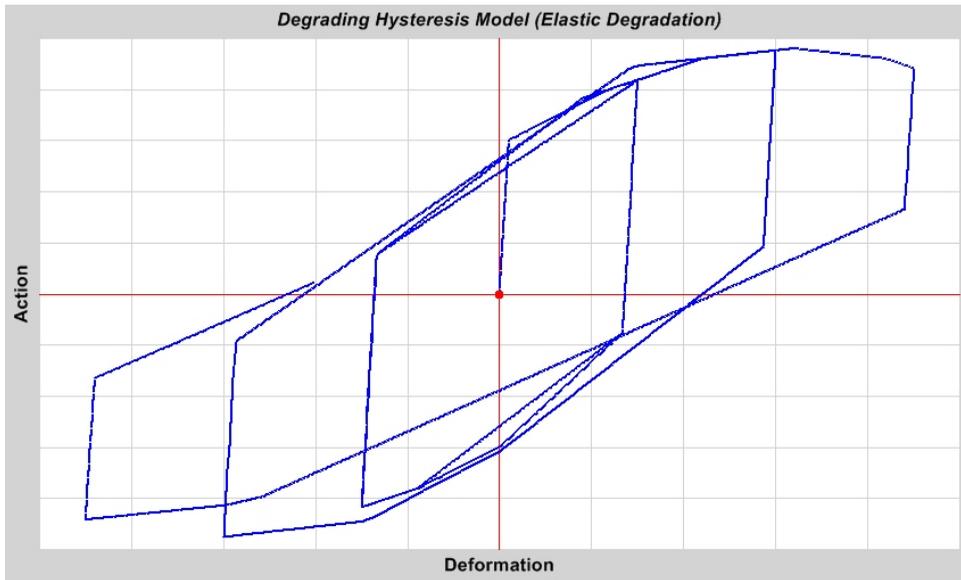


Figure 17
*Degrading Hysteresis Model under Increasing Cyclic Load
 Exhibiting Elastic Degradation ($s = 0.0$)*

The energy factors represent the area of a degraded hysteresis loop divided by the energy of the non-degraded loop, such as for the kinematic model. For example, an energy factor of 0.3 means that a full cycle of deformation would only dissipate 30% of the energy that the non-degraded material would. The energy factors must satisfy $1.0 \geq f_0 \geq f_1 \geq f_2 \geq 0.0$. The deformation levels must satisfy $1.0 < x_1 < x_2$.

All weighting factors may take any value from 0.0 to 1.0, inclusive. Because the accumulated plastic deformation is constantly increasing, it is recommended that the weighting factor a generally be small or zero.

For each increment of deformation:

- The absolute maximum positive and negative plastic deformations that have occurred up to this point in the analysis are determined, $d_{\max \text{ pos}}$ and $d_{\max \text{ neg}}$, as well as the accumulated plastic deformation, d_{acc} .
- A positive plastic deformation level is calculated as

$$d_{\text{pos}} = a d_{\text{acc}} + (1 - a) d_{\max \text{ pos}}$$

where a is the accumulated weighting factor for positive deformation.

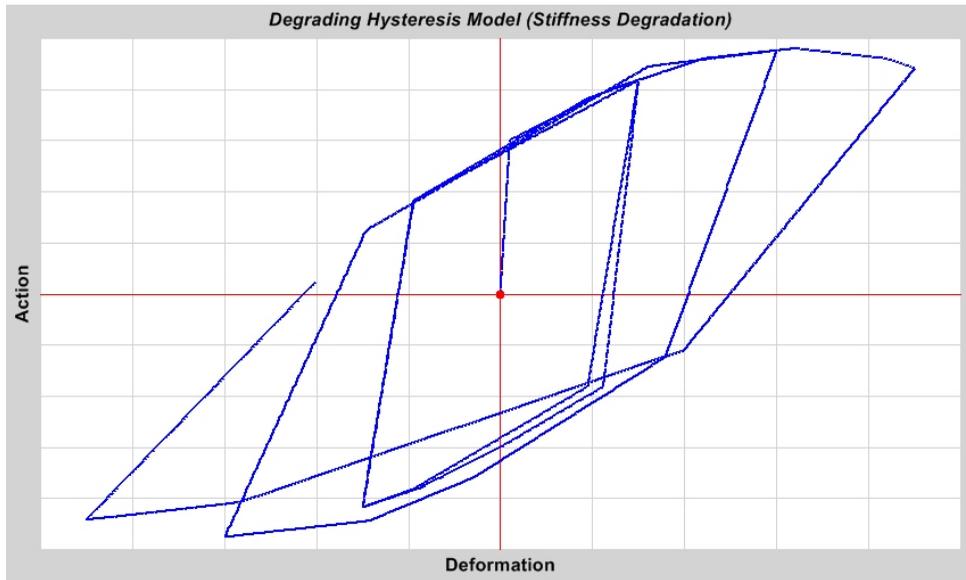


Figure 18
Degrading Hysteresis Model under Increasing Cyclic Load
Exhibiting Stiffness Degradation ($s = 1.0$)

- Comparing d_{pos} with the positive deformation levels $d1$ and $d2$, obtained by multiplying $x1$ and $x2$ with the positive yield deformation, an energy factor f_{pos} can be determined by interpolation. If $d_{pos} > d2$, then $f_{pos} = f2$.
- Following the same approach, the energy factor for negative deformation, f_{neg} , is computed using the corresponding parameters for negative deformation.
- The larger of these two energy factor is called f_{max} , and the smaller is f_{min} . The final energy factor is computed as

$$f = wf_{max} + (1 - w)f_{min}$$

In the most common case, $w = 0$ and $f = f_{min}$.

Degradation does not occur during monotonic loading. However, upon load reversal, the curve for unloading and reverse loading is modified according to the energy factor computed for the last deformation increment. This is done by squeezing, or flattening, the curve toward the diagonal line that connects the two points of maximum positive and negative deformation.

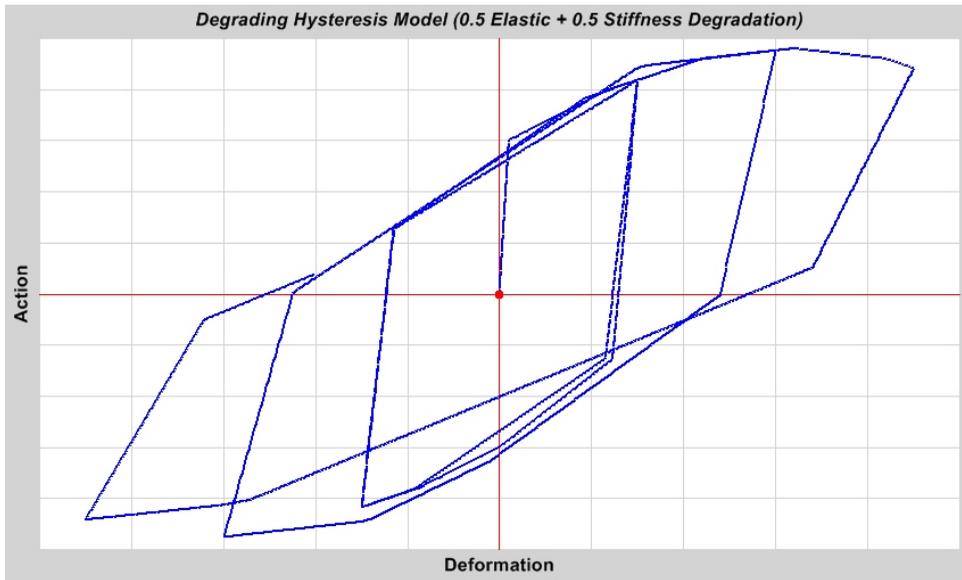


Figure 19
*Degrading Hysteresis Model under Increasing Cyclic Load
Exhibiting Combined Degradation ($s = 0.5$)*

This squeezing is scaled to achieve the desired decrease in energy dissipation. The scaling can occur in two directions:

- Parallel to the elastic unloading line, called elastic degradation
- Parallel to the horizontal axis, called stiffness degradation

The amount of scaling in each direction is controlled by the stiffness degradation weighting parameter, s . For $s = 0.0$, all degradation is of elastic type. For $s = 1.0$, all degradation is of stiffness type. For intermediate values, the degradation is apportioned accordingly.

While the deformation and individual energy levels are computed separately for the positive and negative directions, the final energy level is a single parameter that affects the shape of the hysteresis loop in both directions.

Note that if all the energy factors are equal to 1.0, this model degenerates to the kinematic hysteresis model.

Figures 17, 18, and 19 (pages 90-92) show the shape of the hysteresis loop for elastic degradation, stiffness degradation, and a mixture with a stiffness degradation

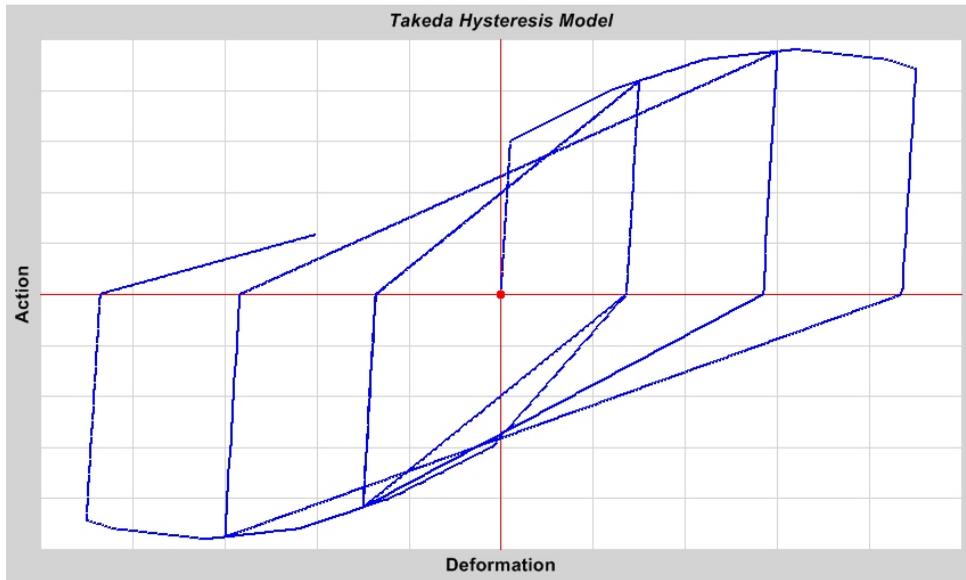


Figure 20
Takeda Hysteresis Model under Increasing Cyclic Load

factor of $s = 0.5$. Each of these three cases dissipates the same amount of energy for a given cycle of loading, and less than the energy dissipated for the equivalent kinematic model shown in Figure 16 (page 88).

Takeda Hysteresis Model

This model is very similar to the kinematic model, but uses a degrading hysteretic loop based on the Takeda model, as described in Takeda, Sozen, and Nielsen (1970). This simple model requires no additional parameters, and is more appropriate for reinforced concrete than for metals. Less energy is dissipated than for the kinematic model.

Unloading is along the elastic segments similar to the kinematic model. When reloading, the curve follows a secant line to the backbone curve for loading in the opposite direction. The target point for this secant is at the maximum deformation that occurred in that direction under previous load cycles. This results in a decreasing amount of energy dissipation with larger deformations. Unloading is along the elastic segments.

This behavior is illustrated in Figure 20 (page 93).

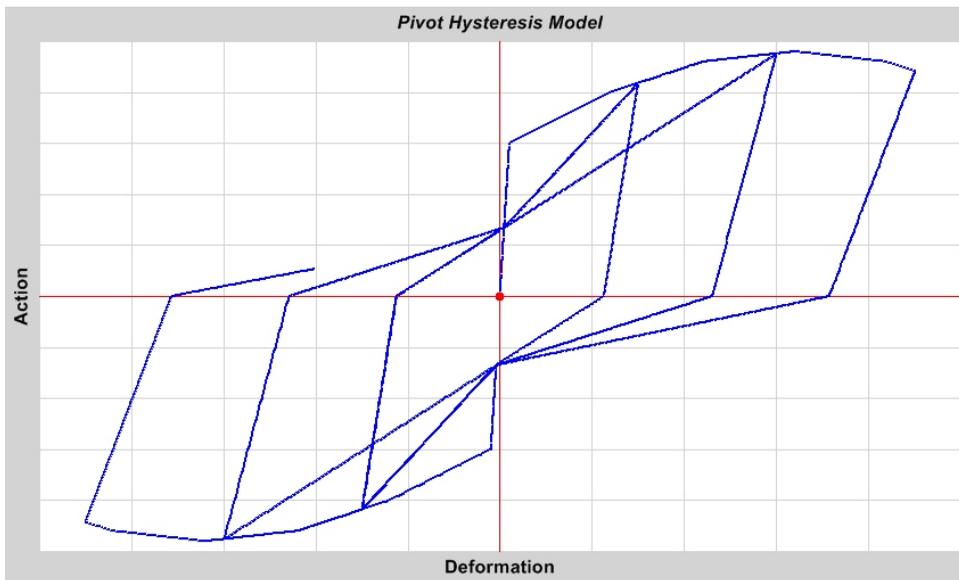


Figure 21
Pivot Hysteresis Model under Increasing Cyclic Load

Pivot Hysteresis Model

This model is similar to the Takeda model, but has additional parameters to control the degrading hysteretic loop. It is particularly well suited for reinforced concrete members, and is based on the observation that unloading and reverse loading tend to be directed toward specific points, called pivot points, in the action-deformation plane. The most common use of this model is for moment-rotation. This model is fully described in Dowell, Seible, and Wilson (1998). This model is not intended for unreinforced concrete. See the separate concrete model below.

The following additional parameters are specified for the Pivot model:

- α_1 , which locates the pivot point for unloading to zero from positive force. Unloading occurs toward a point on the extension of the positive elastic line, but at a negative force value of α_1 times the positive yield force.
- α_2 , which locates the pivot point for unloading to zero from negative force. Unloading occurs toward a point on the extension of the negative elastic line, but at a positive force value of α_2 times the negative yield force.
- β_1 , which locates the pivot point for reverse loading from zero toward positive force. Reloading occurs toward a point on the positive elastic line at a force

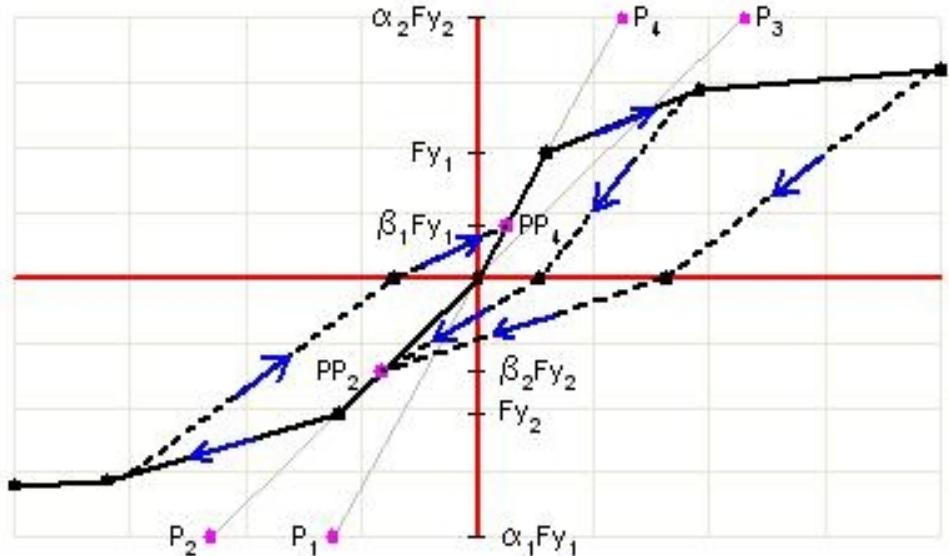


Figure 22
Pivot Hysteresis Model Parameters

value of α_1 times the positive yield force, where $0.0 < \alpha_1 < 1.0$. Beyond that point, loading occurs along the secant to the point of maximum previous positive deformation on the backbone curve.

- α_2 , which locates the pivot point for reverse loading from zero toward negative force. Reloading occurs toward a point on the negative elastic line at a force value of α_2 times the negative yield force, where $0.0 < \alpha_2 < 1.0$. Beyond that point, loading occurs along the secant to the point of maximum previous negative deformation on the backbone curve.
- β , which determines the amount of degradation of the elastic slopes after plastic deformation, where $0.0 < \beta < 1.0$

These parameters and the behavior are illustrated in Figures 21 and 22 (pages 94 and 95).

Concrete Hysteresis Model

This model is intended for unreinforced concrete and similar materials, and is the default model for concrete and masonry materials in the program. Tension and

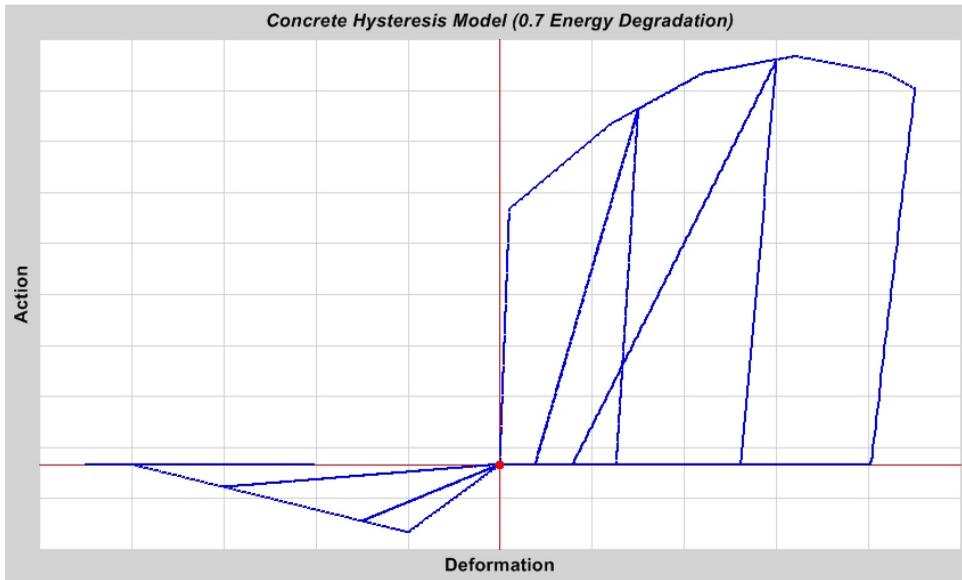


Figure 23
*Concrete Hysteresis Model under Increasing Cyclic Load
 with Compression as Positive and Energy Factor $f = 0.7$*

compression behavior are independent and behave differently. The force-deformation (stress-strain) curve is used to determine the sign of compression, which can be positive or negative. The point having the largest absolute value of stress or force is considered to be in compression, so that the sign of compression can be either positive or negative. Likewise, the concrete model can also be used to represent a tension-only material whose behavior is similar to concrete in compression.

This model is primarily intended for axial behavior, but can be applied to any degree of freedom. Reinforced concrete is better modeled using the pivot, degrading, or Takeda models.

A non-zero force-deformation curve should always be defined for compression. The force-deformation curve for tension may be all zero, or it may be non-zero provided that the maximum force value is of smaller magnitude than that for the compression side.

A single parameter, the energy degradation factor f , is specified for this model. This value should satisfy $0.0 \leq f \leq 1.0$. A value of $f = 0.0$ is equivalent to a clean gap when unloading from compression and dissipates the least amount of energy. A

value of $f = 1.0$ dissipates the most energy and could be caused by rubble filling the gap when unloading from compression.

Compression behavior is modeled as follows:

- Initial loading is along the backbone curve
- Unloading to zero occurs along a line nearly parallel to the compression elastic line. The line is actually directed to a pivot point on the extension of the compressive elastic line, located so that the unloading slope at maximum compressive force has half the stiffness of the elastic loading line.
- At zero force, reverse loading toward tension occurs at zero force.
- Subsequent loading in compression occurs along the previous unloading line if the energy factor $f = 0.0$, and along the secant from the origin to the point of maximum previous compressive deformation if the energy factor is 1.0. An intermediate secant from the horizontal axis is used for other values of f .

Tension behavior, if non-zero, is modeled as follows:

- Initial loading is along the backbone curve
- Unloading occurs along a secant line to the origin.
- Subsequent loading occurs along the unloading secant from the origin to the point of maximum previous tensile deformation.

See Figure 23 (page 96) for an example of this behavior for an energy degradation factor of $f = 0.7$.

BRB Hardening Hysteresis Model

This model is similar to the kinematic model, but accounts for the increasing strength with plastic deformation that is typical of bucking-restrained braces, causing the backbone curve, and hence the hysteresis loop, to progressively grow in size. It is intended primarily for use with axial behavior, but can be applied to any degree of freedom.

Two measures are used for plastic deformation:

- Maximum plastic deformation in each the positive and negative directions
- Accumulated plastic deformation, which is the absolute sum of each increment of positive or negative plastic deformation. Plastic deformation is that which does not occur on the two elastic segment of the force-deformation curve

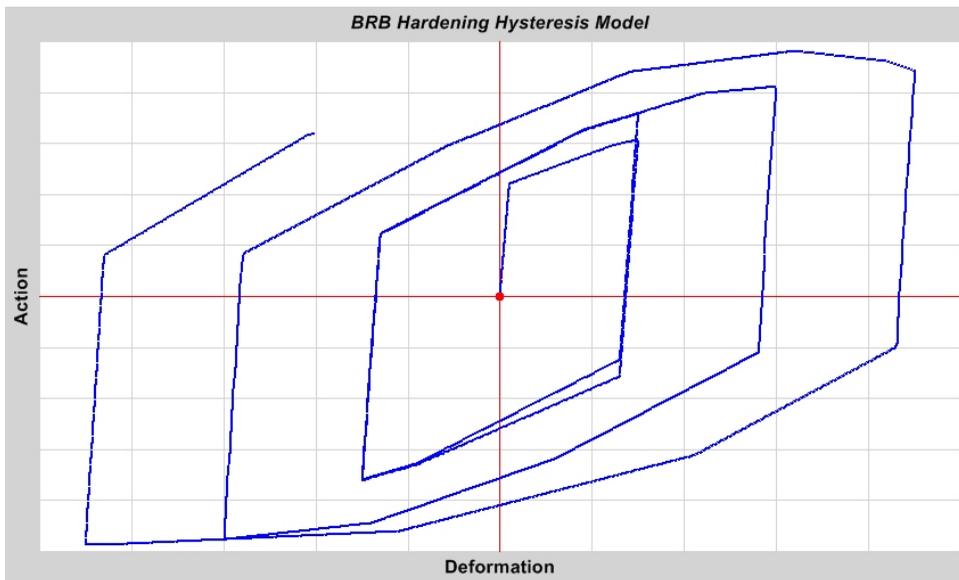


Figure 24
*BRB Hardening Hysteresis Model under Increasing Cyclic Load
 with Hardening Factor $h = 1.5$*

Accumulated plastic deformation can occur under cyclic loading of constant amplitude.

For this model, the following parameters are required:

- Separately for tension (positive) and compression (negative) deformations
 - Hardening factor at maximum deformation, h , where $h \geq 1.0$.
 - Maximum plastic deformation level at full hardening, x_2 , as a multiple of yield deformation, where $x_2 > 1.0$
 - Maximum accumulated plastic deformation level at full hardening, x_4 , as a multiple of yield deformation, where $x_4 > 1.0$
 - Accumulated deformation weighting factor, a , where $0.0 \leq a \leq 1.0$.

The hardening factors scale the size of the backbone curve and hysteresis loop in the action (stress/force/moment) direction. Because the accumulated plastic deformation is constantly increasing, it is recommended that the weighting factor a generally be small or zero.

For each increment of deformation:

- The absolute maximum positive and negative plastic deformations that have occurred up to this point in the analysis are determined, $d_{\max pos}$ and $d_{\max neg}$, as well as the accumulated plastic deformation, d_{acc} .
- Comparing $d_{\max pos}$ with the positive deformation level $d2$, obtained by multiplying $x2$ by the positive yield deformation, a hardening factor $h_{\max pos}$ can be determined by interpolation. If $d_{\max pos} > d2$, then $h_{\max pos} = h$.
- Comparing d_{acc} with the positive deformation level $d4$, obtained by multiplying $x4$ by the positive yield deformation, a hardening factor $h_{acc pos}$ can be determined by interpolation. If $d_{acc} > d4$, then $h_{acc pos} = h$.
- The net hardening factor due to positive deformation, h_{pos} , is computed as

$$h_{pos} = ah_{acc pos} + (1 - a)h_{\max pos}$$

- Following the same approach, the hardening factor due to negative deformation, h_{neg} , is computed using the corresponding parameters for negative deformation.

Degradation does not occur during monotonic loading. However, upon load reversal, the curve for unloading and reverse loading in the opposite direction is modified according to the hardening factor computed for the last deformation increment. This is done by scaling the action values in that direction, including the backbone curve for further loading.

Important! Positive deformation and the corresponding hardening parameters only affect the negative strength, and vice versa.

Note that if the hardening factor is equal to 1.0, this model degenerates to the kinematic hysteresis model.

This behavior is illustrated in Figure 24 (page 98).

Isotropic Hysteresis Model

This model is, in a sense, the opposite of the kinematic model. Plastic deformation in one direction “pushes” the curve for the other direction away from it, so that both directions increase in strength simultaneously. Unlike the BRB hardening model, the backbone curve itself does not increase in strength, only the unloading and reverse loading behavior. Matching pairs of points are linked. No additional parameters are required for this model.

Unloading and reverse loading occur along a path parallel to the elastic line until the magnitude of the action in the reverse direction equals that of backbone curve at

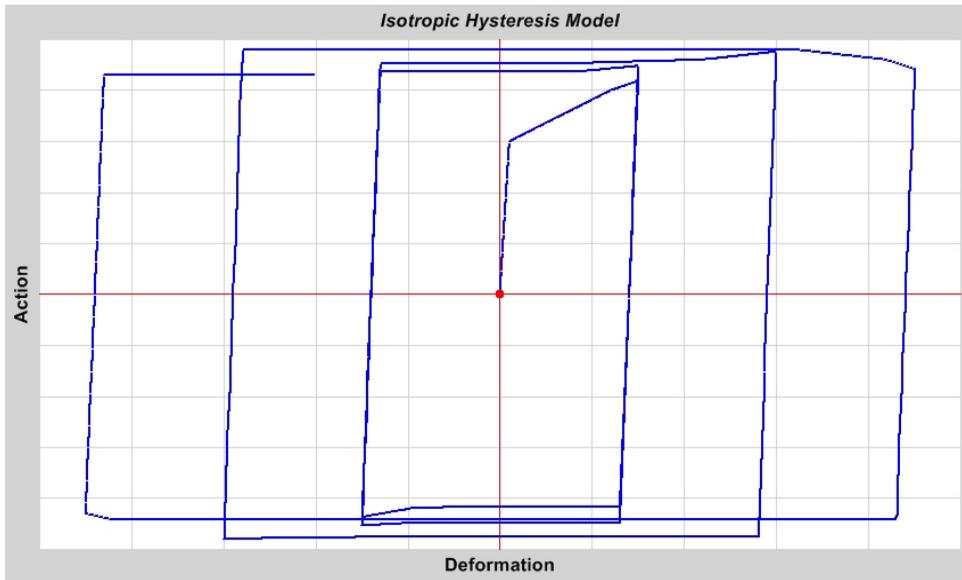


Figure 25
Isotropic Hysteresis Model under Increasing Cyclic Load

the same amount of deformation in the reverse direction, and then continues along a horizontal secant to the backbone curve.

When you define the points on the multi-linear curve, you should be aware that symmetrical pairs of points will be linked, even if the curve is not symmetrical. This gives you some control over the shape of the hysteretic loop.

This model dissipates the most energy of all the models. This behavior is illustrated in Figure 25 (page 100).

Modified Darwin-Pecknold Concrete Model

A two-dimensional nonlinear concrete material model is available for use in the layered shell. This model is based on the Darwin-Pecknold model, with consideration of Vecchio-Collins behavior. This model represents the concrete compression, cracking, and shear behavior under both monotonic and cyclic loading, and considers the stress-strain components ϵ_{11} , ϵ_{22} , and ϵ_{33} . A state of plane stress is assumed.

The direction of cracking can change during the loading history, and the shear strength is affected by the tension strain in the material. The axial stress-strain stress-strain curve specified for the material is simplified to account for initial stiffness, yielding, ultimate plateau, and strength loss due to crushing. Zero tensile strength is assumed.

Hysteresis is governed by the concrete hysteresis model described in the previous topic, with the energy dissipation factor $f = 0$.

The layered shell allows this material to be used for membrane and/or flexural behavior and to be combined with steel reinforcement placed in arbitrary directions and locations. Transverse (out-of-plane) shear is assumed to be elastic and isotropic using the shear stiffness G_{13} for both 13 13 and 23 23 behavior.

See separate Technical Note “Modified Darwin-Pecknold 2-D Reinforced Concrete Material Model” for more information, available using the command Help > Documentation.

Time-dependent Properties

For any material having a design type of concrete or tendon, you may specify time dependent material properties that are used for creep, shrinkage, and aging effects during a staged-construction analysis.

For more information, see Topic “Staged Construction” (page 439) in Chapter “Nonlinear Static Analysis.”

Properties

For concrete-type materials, you may specify:

- Aging parameters that determine the change in modulus of elasticity with age
- Shrinkage parameters that determine the decrease in direct strains with time
- Creep parameters that determine the change in strain with time under the action of stress

For tendon-type materials, relaxation behavior may be specified that determines the change in strain with time under the action of stress, similar to creep.

Currently these behaviors can be specified for the CEB-FIP 1990 code (Comite Euro-International Du Beton, 1993) for concrete and tendon materials, and for con-

crete materials using the following codes: CEB-FIP 2010, ACI 290R-92, and user-specified curves.

Time-Integration Control

For each material, you have the option to model the creep behavior by full integration or by using a Dirichlet series approximation.

With full integration, each increment of stress during the analysis becomes part of the memory of the material. This leads to accurate results, but for long analyses with many stress increments, this requires computer storage and execution time that both increase as the square of the number of increments. For larger problems, this can make solution impractical.

Using the Dirichlet series approximation (Ketchum, 1986), you can choose a fixed number of series terms that are to be stored. Each term is modified by the stress increments, but the number of terms does not change during the analysis. This means the storage and execution time increase linearly with the number of stress increments. Each term in the Dirichlet series can be thought of as a spring and dashpot system with a characteristic relaxation time. The program automatically chooses these spring-dashpot systems based on the number of terms you request. You should try different numbers of terms and check the analysis results to make sure that your choice is adequate.

It is recommended that you work with a smaller problem that is representative of your larger model, and compare various numbers of series terms with the full integration solution to determine the appropriate series approximation to use.

Design-Type

You may specify a **design-type** for each Material that indicates how it is to be treated for design by the SAP2000, ETABS, SAFE, or CSiBridge graphical user interface. The available design types are:

- Steel: Frame elements made of this material will be designed according to steel design codes
- Concrete: Frame elements made of this material will be designed according to concrete design codes
- Aluminum: Frame elements made of this material will be designed according to aluminum design codes

- Cold-formed: Frame elements made of this material will be designed according to cold-formed steel design codes
- None: Frame elements made of this material will not be designed

When you choose a design type, additional material properties may be specified that are used only for design; they do not affect the analysis. Consult the on-line help and design documentation for further information on these design properties.

The Frame Element

The **Frame** element is a very powerful element that can be used to model beams, columns, braces, and trusses in planar and three-dimensional structures. Nonlinear material behavior is available through the use of Frame Hinges.

Basic Topics for All Users

- Overview
- Joint Connectivity
- Degrees of Freedom
- Local Coordinate System
- Section Properties
- Insertion Point
- End Offsets
- End Releases
- Mass
- Self-Weight Load
- Concentrated Span Load
- Distributed Span Load

- Internal Force Output
- Stress Output

Advanced Topics

- Advanced Local Coordinate System
- Property Modifiers
- Nonlinear Properties
- Gravity Load
- Temperature Load
- Strain and Deformation Load
- Target-Force Load

Overview

The Frame element uses a general, three-dimensional, beam-column formulation which includes the effects of biaxial bending, torsion, axial deformation, and biaxial shear deformations. See Bathe and Wilson (1976).

Structures that can be modeled with this element include:

- Three-dimensional frames
- Three-dimensional trusses
- Planar frames
- Planar grillages
- Planar trusses
- Cables

A Frame element is modeled as a straight line connecting two points. In the graphical user interface, you can divide curved objects into multiple straight objects, subject to your specification.

Each element has its own local coordinate system for defining section properties and loads, and for interpreting output.

The element may be prismatic or non-prismatic. The non-prismatic formulation allows the element length to be divided into any number of segments over which properties may vary. The variation of the bending stiffness may be linear, para-

bolic, or cubic over each segment of length. The axial, shear, torsional, mass, and weight properties all vary linearly over each segment.

Insertion points and end offsets are available to account for the finite size of beam and column intersections. The end offsets may be made partially or fully rigid to model the stiffening effect that can occur when the ends of an element are embedded in beam and column intersections. End releases are also available to model different fixity conditions at the ends of the element.

Each Frame element may be loaded by gravity (in any direction), multiple concentrated loads, multiple distributed loads, strain and deformation loads, and loads due to temperature change.

Target-force loading is available that iteratively applies deformation load to the element to achieve a desired axial force.

Element internal forces are produced at the ends of each element and at a user-specified number of equally-spaced output stations along the length of the element.

Cable behavior is usually best modeled using the catenary Cable element (page 165). However, there are certain cases where using Frame elements is necessary. This can be achieved by adding appropriate features to a Frame element. You can release the moments at the ends of the elements, although we recommend that you retain small, realistic bending stiffness instead. You can also add nonlinear behavior as needed, such as the no-compression property, tension stiffening (p-delta effects), and large deflections. These features require nonlinear analysis.

Joint Connectivity

A Frame element is represented by a straight line connecting two joints, **I** and **j**, unless modified by insertion points as described below. The two joints must not share the same location in space. The two ends of the element are denoted End I and End J, respectively.

Insertion Points

Sometimes the neutral axis of the element cannot be conveniently located by joints that connect to other elements in the structure. You have the option to specify **insertion points** that locate the element with respect to the joints. The insertion points consists of a **cardinal point** specified for the section, plus independent **joint offsets** specified at each end of the element. By default the cardinal point is the centroid of the section and the joints offsets are zero.

The two ends of the neutral axis, considering the coordinates of joints **I** and **j** plus the insertion points, must not be coincident. It is generally recommended that the offsets due to the insertion points be perpendicular to the axis of the element, although this is not required.

For more information on the insertion points, including how they affect the local coordinate system of the element, see Topic “Insertion Points” (page 125).

Degrees of Freedom

The Frame element activates all six degrees of freedom at both of its connected joints. If you want to model truss or cable elements that do not transmit moments at the ends, you may either:

- Set the geometric Section properties **j**, **i33**, and **i22** all to zero (**a** is non-zero; **as2** and **as3** are arbitrary), or
- Release both bending rotations, R2 and R3, at both ends and release the torsional rotation, R1, at either end

In either case, the joint offsets and the end offsets must both be zero to avoid moments at the ends.

For more information:

- See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom.”
- See Topic “Section Properties” (page 114) in this Chapter.
- See Topic “End Offsets” (page 127) in this Chapter.
- See Topic “End Releases” (page 131) in this Chapter.

Local Coordinate System

Each Frame element has its own **element local coordinate system** used to define section properties, loads and output. The axes of this local system are denoted 1, 2 and 3. The first axis is directed along the length of the element at its centroid; the remaining two axes lie in the plane perpendicular to the element with an orientation that you specify.

It is important that you clearly understand the definition of the element local 1-2-3 coordinate system and its relationship to the global X-Y-Z coordinate system. Both

systems are right-handed coordinate systems. It is up to you to define local systems which simplify data input and interpretation of results.

In most structures the definition of the element local coordinate system is extremely simple. The methods provided, however, provide sufficient power and flexibility to describe the orientation of Frame elements in the most complicated situations.

Local axes are first computed for the element without considering the insertion points, i.e., as if the neutral axis connects the two joints. These are called the nominal local axes. If the insertion points shift the neutral axis by a different amount at the two ends, the local axes are then transformed by projecting them onto the neutral axis to determine the actual local coordinate system used for analysis.

The discussion below considers the calculation of the nominal local axes using the joints. The transformation for the insertion points, if needed, is discussed later in Topic “Insertion Points”.

The simplest method for computing the element local coordinate system, using the **default orientation** and the **Frame element coordinate angle**, is described in this topic. Additional methods for defining the Frame element local coordinate system are described in the next topic.

For more information:

- See Chapter “Coordinate Systems” (page 11) for a description of the concepts and terminology used in this topic.
- See Topic “Advanced Local Coordinate System” (page 110) in this Chapter.
- See Topic “Insertion Points” (page 125) in this Chapter.

Longitudinal Axis I

The local axis 1 is always the longitudinal axis of the element, the positive direction being directed from End I to End J. This axis is always located at the centroid of the cross section, and connects joint **I** to joint **j**.

Default Orientation

The default orientation of the local 2 and 3 axes is determined by the relationship between the local 1 axis and the global Z axis. The local 1 axis is directed along the line between the joints **I** and **j** without considering any offsets:

- The local 1-2 plane is taken to be vertical, i.e., parallel to the Z axis

- The local 2 axis is taken to have an upward (+Z) sense unless the element is vertical, in which case the local 2 axis is taken to be horizontal along the global +X direction
- The local 3 axis is horizontal, i.e., it lies in the X-Y plane

An element is considered to be vertical if the sine of the angle between the local 1 axis and the Z axis is less than 10^{-3} .

The local 2 axis makes the same angle with the vertical axis as the local 1 axis makes with the horizontal plane. This means that the local 2 axis points vertically upward for horizontal elements.

Coordinate Angle

The Frame element coordinate angle, **ang**, is used to define element orientations that are different from the default orientation. It is the angle through which the local 2 and 3 axes are rotated about the positive local 1 axis from the default orientation. The rotation for a positive value of **ang** appears counterclockwise when the local +1 axis is pointing toward you.

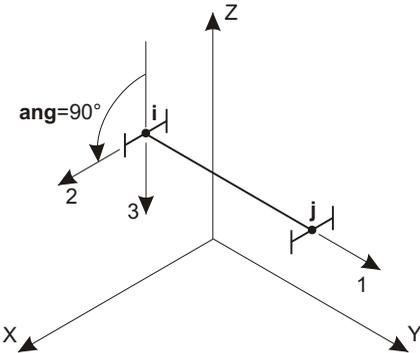
For vertical elements, **ang** is the angle between the local 2 axis and the horizontal +X axis. Otherwise, **ang** is the angle between the local 2 axis and the vertical plane containing the local 1 axis. See Figure 26 (page 111) for examples.

Advanced Local Coordinate System

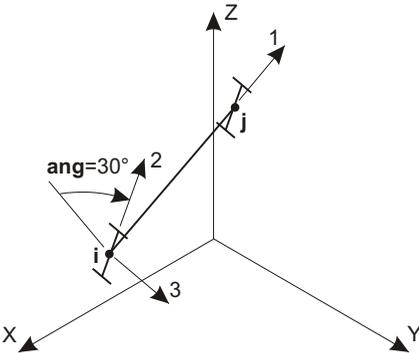
By default, the element local coordinate system is defined using the element coordinate angle measured with respect to the global +Z and +X directions, as described in the previous topic. In certain modeling situations it may be useful to have more control over the specification of the local coordinate system.

This topic describes how to define the orientation of the transverse local 2 and 3 axes with respect to an arbitrary reference vector when the element coordinate angle, **ang**, is zero. If **ang** is different from zero, it is the angle through which the local 2 and 3 axes are rotated about the positive local 1 axis from the orientation determined by the reference vector. The local 1 axis is always directed from end I to end J of the element.

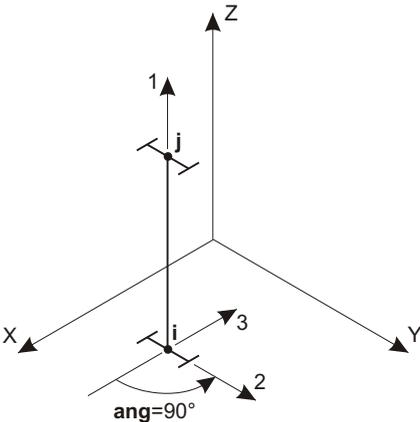
The discussion below considers the calculation of the nominal local axes using the joints. The transformation for the insertion points, if needed, is discussed later in Topic “Insertion Points”.



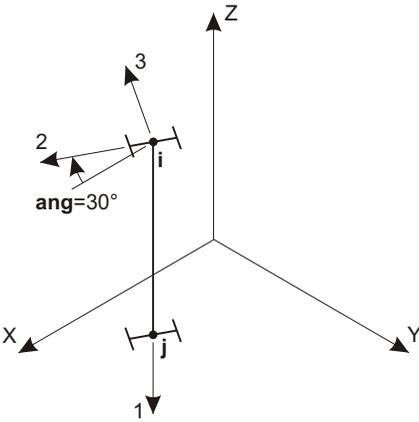
Local 1 Axis is Parallel to +Y Axis
Local 2 Axis is Rotated 90° from Z-1 Plane



Local 1 Axis is Not Parallel to X, Y, or Z Axes
Local 2 Axis is Rotated 30° from Z-1 Plane



Local 1 Axis is Parallel to +Z Axis
Local 2 Axis is Rotated 90° from X-1 Plane



Local 1 Axis is Parallel to -Z Axis
Local 2 Axis is Rotated 30° from X-1 Plane

Figure 26

The Frame Element Coordinate Angle with Respect to the Default Orientation

For more information:

- See Chapter “Coordinate Systems” (page 11) for a description of the concepts and terminology used in this topic.

- See Topic “Local Coordinate System” (page 108) in this Chapter.
- See Topic “Insertion Points” (page 125) in this Chapter.

Reference Vector

To define the transverse local axes 2 and 3, you specify a **reference vector** that is parallel to the desired 1-2 or 1-3 plane. The reference vector must have a positive projection upon the corresponding transverse local axis (2 or 3, respectively). This means that the positive direction of the reference vector must make an angle of less than 90° with the positive direction of the desired transverse axis.

To define the reference vector, you must first specify or use the default values for:

- A primary coordinate direction **pldirp** (the default is +Z)
- A secondary coordinate direction **pldirs** (the default is +X). Directions **pldirs** and **pldirp** should not be parallel to each other unless you are sure that they are not parallel to local axis 1
- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system)
- The local plane, **local**, to be determined by the reference vector (the default is 12, indicating plane 1-2)

You may optionally specify:

- A pair of joints, **plveca** and **plvecb** (the default for each is zero, indicating the center of the element). If both are zero, this option is not used

For each element, the reference vector is determined as follows:

1. A vector is found from joint **plveca** to joint **plvecb**. If this vector is of finite length and is not parallel to local axis 1, it is used as the reference vector \mathbf{V}_p
2. Otherwise, the primary coordinate direction **pldirp** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis 1, it is used as the reference vector \mathbf{V}_p
3. Otherwise, the secondary coordinate direction **pldirs** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis 1, it is used as the reference vector \mathbf{V}_p
4. Otherwise, the method fails and the analysis terminates. This will never happen if **pldirp** is not parallel to **pldirs**

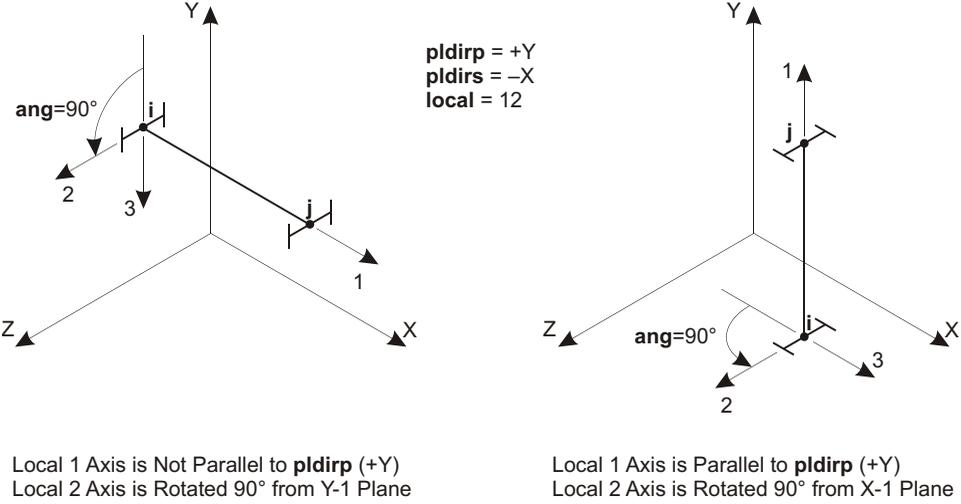


Figure 27
The Frame Element Coordinate Angle with Respect to Coordinate Directions

A vector is considered to be parallel to local axis 1 if the sine of the angle between them is less than 10^{-3} .

The use of the Frame element coordinate angle in conjunction with coordinate directions that define the reference vector is illustrated in Figure 27 (page 113). The use of joints to define the reference vector is shown in Figure 28 (page 114).

Determining Transverse Axes 2 and 3

The program uses vector cross products to determine the transverse axes 2 and 3 once the reference vector has been specified. The three axes are represented by the three unit vectors V_1 , V_2 and V_3 , respectively. The vectors satisfy the cross-product relationship:

$$V_1 \times V_2 = V_3$$

The transverse axes 2 and 3 are defined as follows:

- If the reference vector is parallel to the 1-2 plane, then:

$$V_3 = V_1 \times V_p \text{ and } V_2 = V_3 \times V_1$$

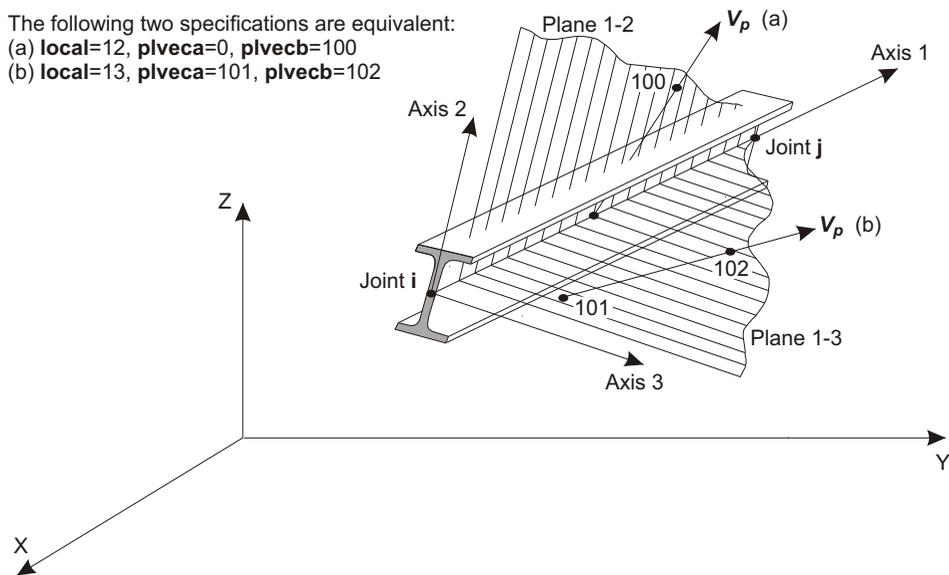


Figure 28
Using Joints to Define the Frame Element Local Coordinate System

- If the reference vector is parallel to the 1-3 plane, then:

$$\begin{matrix} \mathbf{V}_2 & \mathbf{V}_p & \mathbf{V}_1 & \text{and} \\ \mathbf{V}_3 & \mathbf{V}_1 & \mathbf{V}_2 \end{matrix}$$

In the common case where the reference vector is perpendicular to axis \mathbf{V}_1 , the transverse axis in the selected plane will be equal to \mathbf{V}_p .

Section Properties

A **Frame Section** is a set of material and geometric properties that describe the cross-section of one or more Frame elements. Sections are defined independently of the Frame elements, and are assigned to the elements.

Section properties are of two basic types:

- Prismatic — all properties are constant along the full element length
- Non-prismatic — the properties may vary along the element length

Non-prismatic Sections are defined by referring to two or more previously defined prismatic Sections.

All of the following subtopics, except the last, describe the definition of prismatic Sections. The last subtopic, “Non-prismatic Sections”, describes how prismatic Sections are used to define non-prismatic Sections.

Local Coordinate System

Section properties are defined with respect to the local coordinate system of a Frame element as follows:

- The 1 direction is along the axis of the element. It is normal to the Section and goes through the intersection of the two neutral axes of the Section.
- The 2 and 3 directions are parallel to the neutral axes of the Section. Usually the 2 direction is taken along the major dimension (depth) of the Section, and the 3 direction along its minor dimension (width), but this is not required.

See Topic “Local Coordinate System” (page 108) in this Chapter for more information.

Material Properties

The material properties for the Section are specified by reference to a previously-defined Material. Isotropic material properties are used, even if the Material selected was defined as orthotropic or anisotropic. The material properties used by the Section are:

- The modulus of elasticity, **e1**, for axial stiffness and bending stiffness
- The shear modulus, **g12**, for torsional stiffness and transverse shear stiffness
- The coefficient of thermal expansion, **a1**, for axial expansion and thermal bending strain
- The mass density, **m**, for computing element mass
- The weight density, **w**, for computing Self-Weight and Gravity Loads

The material properties **e1**, **g12**, and **a1** are all obtained at the material temperature of each individual Frame element, and hence may not be unique for a given Section.

See Chapter “Material Properties” (page 69) for more information.

Geometric Properties and Section Stiffnesses

Six basic geometric properties are used, together with the material properties, to generate the stiffnesses of the Section. These are:

- The cross-sectional area, **a**. The axial stiffness of the Section is given by **a e1**;
- The moment of inertia, **i33**, about the 3 axis for bending in the 1-2 plane, and the moment of inertia, **i22**, about the 2 axis for bending in the 1-3 plane. The corresponding bending stiffnesses of the Section are given by **i33 e1** and **i22 e1**;
- The torsional constant, **j**. The torsional stiffness of the Section is given by **j g12**. *Note that the torsional constant is not the same as the polar moment of inertia, except for circular shapes.* See Roark and Young (1975) or Cook and Young (1985) for more information.
- The shear areas, **as2** and **as3**, for transverse shear in the 1-2 and 1-3 planes, respectively. The corresponding transverse shear stiffnesses of the Section are given by **as2 g12** and **as3 g12**. Formulae for calculating the shear areas of typical sections are given in Figure 29 (page 117).

Setting **a**, **j**, **i33**, or **i22** to zero causes the corresponding section stiffness to be zero. For example, a truss member can be modeled by setting **j = i33 = i22 = 0**, and a planar frame member in the 1-2 plane can be modeled by setting **j = i22 = 0**.

Setting **as2** or **as3** to zero causes the corresponding transverse shear *deformation* to be zero. In effect, a zero shear area is interpreted as being infinite. The transverse shear stiffness is ignored if the corresponding bending stiffness is zero.

Shape Type

For each Section, the six geometric properties (**a**, **j**, **i33**, **i22**, **as2** and **as3**) may be specified directly, computed from specified Section dimensions, or read from a specified property database file. This is determined by the shape type, **shape**, specified by the user:

- If **shape**=GENERAL (general section), the six geometric properties must be explicitly specified
- If **shape**=RECTANGLE, PIPE, BOX/TUBE, I/WIDE FLANGE, or one of several others offered by the program, the six geometric properties are automatically calculated from specified Section dimensions as described in “Automatic Section Property Calculation” below, or obtained from a specified property database file. See “Section Property Database Files” below.

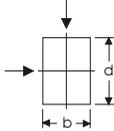
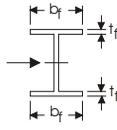
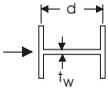
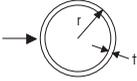
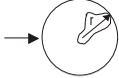
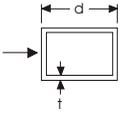
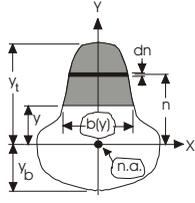
Section	Description	Effective Shear Area
	Rectangular Section Shear Forces parallel to the b or d directions	$\frac{5}{6} bd$
	Wide Flange Section Shear Forces parallel to flange	$\frac{5}{3} t_f b_f$
	Wide Flange Section Shear Forces parallel to web	$t_w d$
	Thin Walled Circular Tube Section Shear Forces from any direction	$(0.9 - 0.4 s) A$ where $s = (r - t) / r$ $A = \pi (2r - t) t$
	Solid Circular Section Shear Forces from any direction	$0.9 \pi r^2$
	Thin Walled Rectangular Tube Section Shear Forces parallel to d-direction	$2 t d$
	General Section Shear Forces parallel to Y-direction $I_x =$ moment of inertia of section about X-X $Q(Y) = \int_y^{Y_t} n b(n) dn$	$\frac{I_x^2}{\int_b^{Y_t} \frac{Q^2(y)}{b(y)} dy}$

Figure 29
Shear Area Formulae

- If **shape=SD SECTION** (Section Designer Section), you can create your own arbitrary Sections using the Section Designer utility within the program, and

the six geometric properties are automatically calculated. See “Section Designer Sections” below.

- If **shape**=NONPRISMATIC, the Section is interpolated along the length of the element from previously defined Sections as described in “Nonprismatic Section” below.

Automatic Section Property Calculation

The six geometric Section properties can be automatically calculated from specified dimensions for the simple shapes shown in Figure 30 (page 119), and for others offered by the program. The required dimensions for each shape are shown in the figure.

Note that the dimension **t3** is the depth of the Section in the 2 direction and contributes primarily to **i33**.

Section Property Database Files

Geometric Section properties may be obtained from one or more Section property database files. Several database files are currently supplied with SAP2000, ETABS, or CSiBridge that provide properties for different regions, codes, and manufacturers. Their format and content may be different for each product. Additional property database files may be created by the user or may be available from other sources.

The geometric properties are stored in the length units specified when the database file was created. These are automatically converted to the appropriate units when used in a model. Different frame sections can be obtained from different database files for the same model.

Section-Designer Sections

Section Designer is a separate utility built into SAP2000, ETABS, and CSiBridge that can be used to create your own frame section properties. You can build sections of arbitrary geometry and combinations of materials. The basic analysis geometric properties (areas, moments of inertia, and torsional constant) are computed and used for analysis. In addition, Section Designer can compute nonlinear frame hinge properties.

For more information, see the on-line help within Section Designer.

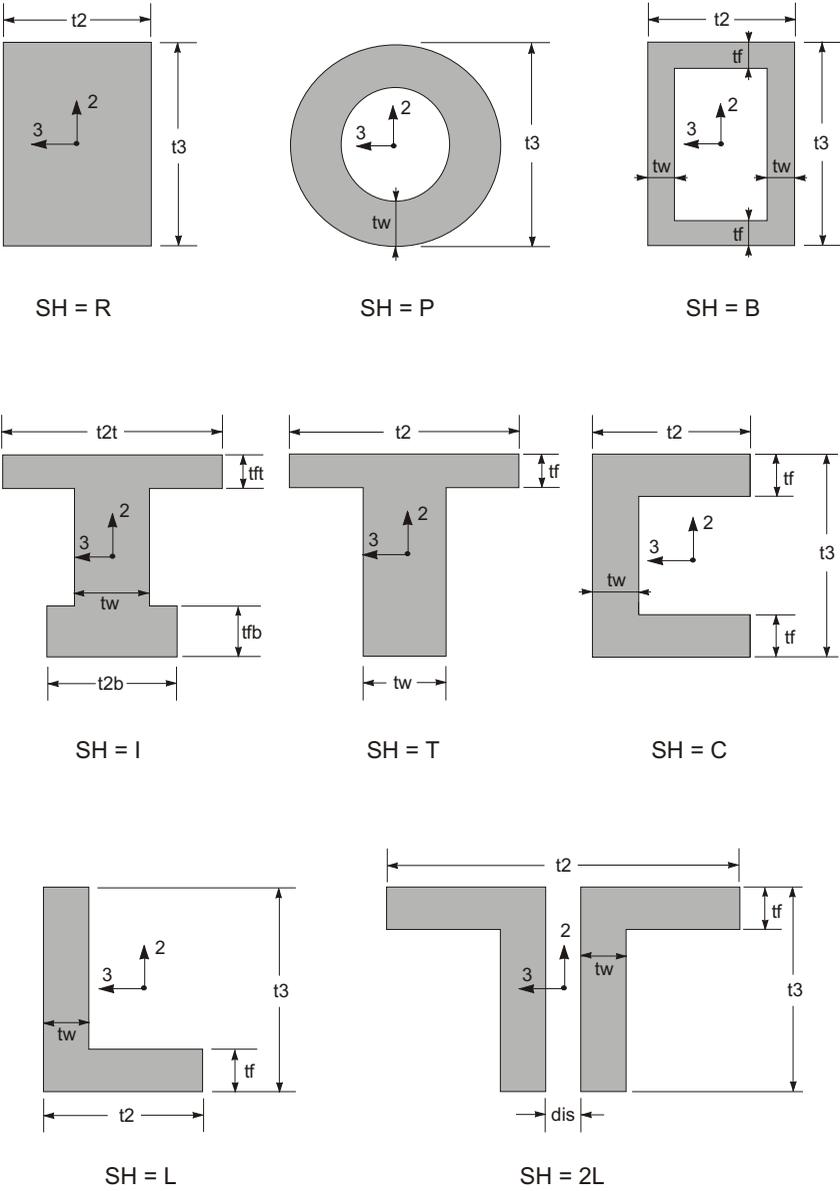


Figure 30
Automatic Section Property Calculation

Additional Mass and Weight

You may specify mass and/or weight for a Section that acts *in addition to* the mass and weight of the material. The additional mass and weight are specified *per unit of length* using the parameters **mpl** and **wpl**, respectively. They could be used, for example, to represent the effects of nonstructural material that is attached to a Frame element.

The additional mass and weight act regardless of the cross-sectional area of the Section. The default values for **mpl** and **wpl** are zero for all shape types.

Non-prismatic Sections

Non-prismatic Sections may be defined for which the properties vary along the element length. You may specify that the element length be divided into any number of segments; these do not need to be of equal length. Most common situations can be modeled using from one to five segments.

The variation of the bending stiffnesses may be linear, parabolic, or cubic over each segment of length. The axial, shear, torsional, mass, and weight properties all vary linearly over each segment. Section properties may change discontinuously from one segment to the next.

See Figure 31 (page 122) for examples of non-prismatic Sections.

Segment Lengths

The length of a non-prismatic segment may be specified as either a variable length, **vl**, or an absolute length, **l**. The default is **vl** = 1.

When a non-prismatic Section is assigned to an element, the actual lengths of each segment for that element are determined as follows:

- The clear length of the element, L_c , is first calculated as the total length minus the end offsets:

$$L_c = L - (\text{ioff} + \text{joff})$$

See Topic “End Offsets” (page 127) in this Chapter for more information.

- If the sum of the absolute lengths of the segments exceeds the clear length, they are scaled down proportionately so that the sum equals the clear length. Otherwise the absolute lengths are used as specified.
- The remaining length (the clear length minus the sum of the absolute lengths) is divided among the segments having variable lengths in the same proportion as

the specified lengths. For example, for two segments with $v1 = 1$ and $v1 = 2$, one third of the remaining length would go to the first segment, and two thirds to the second segment.

Starting and Ending Sections

The properties for a segment are defined by specifying:

- The label, **seci**, of a previously defined prismatic Section that defines the properties at the start of the segment, i.e., at the end closest to joint **I**.
- The label, **secj**, of a previously defined prismatic Section that defines the properties at the end of the segment, i.e., at the end closest to joint **j**. The starting and ending Sections may be the same if the properties are constant over the length of the segment.

The Material would normally be the same for both the starting and ending Sections and only the geometric properties would differ, but this is not required.

Variation of Properties

Non-prismatic Section properties are interpolated along the length of each segment from the values at the two ends.

The variation of the bending stiffnesses, **i33 e1** and **i22 e1**, are defined by specifying the parameters **eivar33** and **eivar22**, respectively. Assign values of 1, 2, or 3 to these parameters to indicate variation along the length that is linear, parabolic, or cubic, respectively.

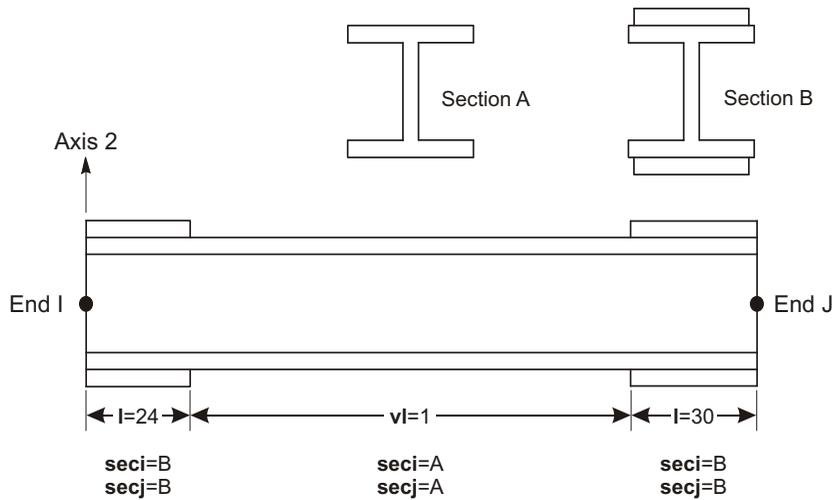
Specifically, the **eivar33**-th root of the bending stiffness in the 1-2 plane:

$$eivar33 \sqrt{i33 e1}$$

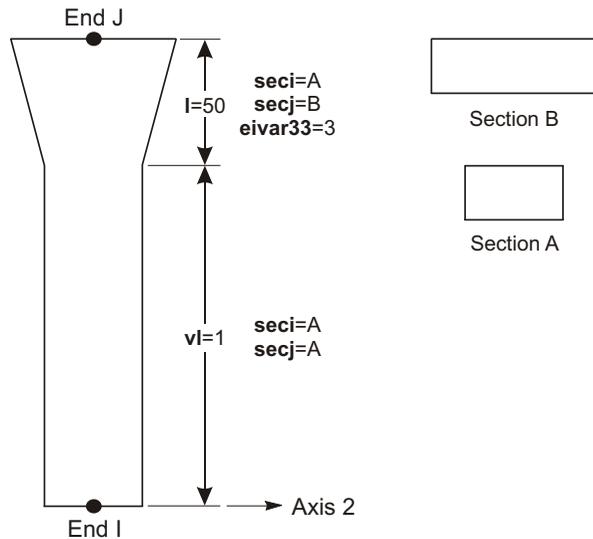
varies linearly along the length. This usually corresponds to a linear variation in one of the Section dimensions. For example, referring to Figure 30 (page 119): a linear variation in **t2** for the rectangular shape would require **eivar33=1**, a linear variation in **t3** for the rectangular shape would require **eivar33=3**, and a linear variation in **t3** for the I-shape would require **eivar33=2**.

The interpolation of the bending stiffness in the 1-2 plane, **i22 e1**, is defined in the same manner by the parameter **eivar22**.

The remaining properties are assumed to vary linearly between the ends of each segment:



Steel Beam with Cover Plates at Ends



Concrete Column with Flare at Top

Figure 31
Examples of Non-prismatic Sections

- Stiffnesses: **a e1, j g12, as2 g12, and as3 g12**
- Mass: **a m + mpl**

- Weight: $a w + wpl$

If a shear area is zero at either end, it is taken to be zero along the full segment, thus eliminating all shear deformation in the corresponding bending plane for that segment.

Advanced Location Parameters

Normally the full variation of a non-prismatic section occurs over the length of a single frame object. When the object is auto-meshed into multiple frame elements, each element will represent a portion of the full non-prismatic length.

Similarly, if you explicitly divide a non-prismatic frame object into multiple frame objects, it is necessary to specify for each frame object what portion of the total nonprismatic variation applies to each object. This is done by assigning to each object:

- The length of the total non-prismatic section. This will be the same for each object derived from a single parent frame object, and will be longer than each derived object.
- The relative location

For example, consider a single non-prismatic frame object of length 8 meters, divided into four equal-length objects. For each derived object, the assigned non-prismatic length should be 8 meters and the relative starting locations should be 0.0, 0.25, 0.5, and 0.75, respectively.

Effect upon End Offsets

Properties vary only along the clear length of the element. Section properties within end offset **ioff** are constant using the starting Section of the first segment. Section properties within end offset **joff** are constant using the ending Section of the last segment.

See Topic “End Offsets” (page 127) in this Chapter for more information.

Property Modifiers

You may specify scale factors to modify the computed section properties. These may be used, for example, to account for cracking of concrete or for other factors not easily described in the geometry and material property values. Individual modifiers are available for the following eight terms:

- The axial stiffness **a e1**
- The shear stiffnesses **as2 g12** and **as3 g12**
- The torsional stiffness **j g12**
- The bending stiffnesses **i33 e1** and **i22 e1**
- The section mass **a m + mpl**
- The section weight **a w + wpl**

You may specify multiplicative factors in two places:

- As part of the definition of the section property
- As an assignment to individual elements.

If modifiers are assigned to an element and also to the section property used by that element, then both sets of factors multiply the section properties. Modifiers cannot be assigned directly to a nonprismatic section property, but any modifiers applied to the sections contributing to the nonprismatic section are used.

When performing steel frame design using the Direct Analysis Method of design code AISC 360-05/IBC2006, further property modifiers may be computed by the design algorithm for the axial and bending stiffnesses. In this case, the computed modifiers are multiplied by those assigned to the element and those specified in the section property used by that element, so that all three sets of factors apply.

Named Property Sets

In addition to directly assigning property modifiers to frame elements, you can apply them to a frame element in a staged-construction Load Case using a Named Property Set of Frame Property Modifiers. A Named Property Set includes the same eight factors above that can be assigned to an element.

When a Named Property Set is applied to an element in a particular stage of a Load Case, it replaces only the values that are assigned to the element or that had been applied in a previous stage; values computed by the Direct Analysis Method of design are also replaced. However, property modifiers specified with the section property remain in force and are not affected by the application of a Named Property Set. The net effect is to use the factors specified in the Named Property Set multiplied by the factors specified in the section property.

When property modifiers are changed in a staged construction Load Case, they do not change the response of the structure up to that stage, but only affect subsequent response. In other words, the effect is incremental. For example, consider a cantile-

ver with only default (unity) property modifiers, and a staged construction case as follows:

- Stage 1: Self-weight load is applied, resulting in a tip deflection of 1.0 and a support moment of 1000.
- Stage 2: Named Property Set “A” is applied that multiplies all stiffnesses by 2.0, and the mass and weight by 1.0. The tip deflection and support moment do not change.
- Stage 3: Self-weight load is applied again (incrementally). The resulting tip deflection is 1.5 and the support moment is 2000. Compared to Stage 1, the same incremental load is applied, but the structure is twice as stiff.
- Stage 4: Named Property Set “B” is applied that multiplies all stiffnesses, as well as the mass and weight, by 2.0. The tip deflection and support moment do not change.
- Stage 5: Self-weight load is applied again (incrementally). The resulting tip deflection is 2.5 and the support moment is 4000. Compared to Stage 1, twice the incremental load is applied, and the structure is twice as stiff

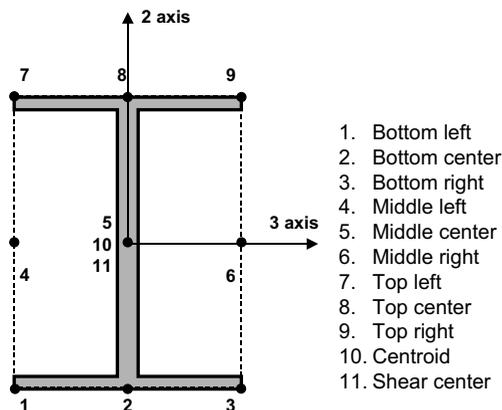
Insertion Points

The local 1 axis of the element runs along the neutral axis of the section, i.e., at the centroid of the section. By default this connects to the joints **I** and **j** at the ends of the element. However, it is often convenient to specify another location on the section, such as the top of a beam or an outside corner of a column, to connect to the joints.

There is a set of pre-defined locations within the section, called **cardinal points**, that can be used for this purpose. The available choices are shown in Figure 32 (page 126). The default location is point 10, the centroid.

You can further offset the cardinal point from the joint by specifying **joint offsets**. The joint offsets together with the cardinal point make up the **insertion point** assignment. The total offset from the joint to the centroid is given as the sum of the joint offset plus the distance from the cardinal point to the centroid.

This feature is useful, as an example, for modeling beams and columns when the beams do not frame into the center of the column. Figure 33 (page 128) shows an elevation and plan view of a common framing arrangement where the exterior beams are offset from the column center lines to be flush with the exterior of the building.



Note: For doubly symmetric members such as this one, cardinal points 5, 10, and 11 are the same.

Figure 32
Frame Cardinal Points

Also shown in this figure are the cardinal points for each member and the joint offset dimensions.

Offsets along the neutral axis of the element are usually specified using **end offsets** rather than insertion points. See topic “End Offsets” (page 127). End offsets are treated as part of the length of the element, have element properties and loads, and may or may not be rigid.

Offsets due to insertion points are external to the element and do not have any mass or loads. Internally the analysis represents the insertion point by a fully rigid constraint that connects the neutral axis to the joints.

Local Axes

The insertion points can interact significantly with the element local coordinate system. As described previously, the nominal local axes are computed for the default insertion points, such that the local 1 axis connects joints **I** and **j**.

The centroids of the section are then located using both the cardinal point and joint offsets. Joint offsets may be specified using the global coordinate system or the element local system. In the latter case, *the nominal local axes are used for this purpose*.

If the neutral axis of the frame section remains parallel to the nominal local 1 axis (the line connecting the two joints), then no further transformation is needed. The element local axes are the same as the nominal axes.

If the neutral axis has changed direction, then the element local coordinate system is computed as follows:

$$\begin{array}{l} \tilde{\mathbf{V}}_3 \quad \tilde{\mathbf{V}}_1 \quad \mathbf{V}_2 \quad \text{and} \\ \tilde{\mathbf{V}}_2 \quad \tilde{\mathbf{V}}_3 \quad \tilde{\mathbf{V}}_1 \end{array}$$

where \mathbf{V}_1 , \mathbf{V}_2 , and \mathbf{V}_3 and the nominal local axes computed previously based on the joints; and $\tilde{\mathbf{V}}_1$, $\tilde{\mathbf{V}}_2$, and $\tilde{\mathbf{V}}_3$ are the transformed local axes used for analysis. If \mathbf{V}_2 is vertical, then $\tilde{\mathbf{V}}_3$ will always be horizontal. Note that the two systems are identical if $\tilde{\mathbf{V}}_1 = \mathbf{V}_1$, the usual case.

The nominal axes are used only for determining the direction of joint offsets. The transformed axes are used for all analysis and design purposes, including loading and results output.

For non-prismatic elements with non-centroidal cardinal points, the local 1 axis may not be straight, and as a consequence the local 2 and 3 axes may change orientation between segments. This can be expected to cause jumps in the axial force, shear, and moments. However, the change will be a small deviation from the axes that would have been calculated for an element with no insertion points.

End Offsets

Frame elements are modeled as line elements connected at points (joints). However, actual structural members have finite cross-sectional dimensions. When two elements, such as a beam and column, are connected at a joint there is some overlap of the cross sections. In many structures the dimensions of the members are large and the length of the overlap can be a significant fraction of the total length of a connecting element.

You may specify two **end offsets** for each element using parameters **ioff** and **joff** corresponding to ends I and J, respectively. End offset **ioff** is the length of overlap for a given element with other connecting elements at joint **I**. It is the distance from the joint to the face of the connection for the given element. A similar definition applies to end offset **joff** at joint **j**. See Figure 34 (page 129).

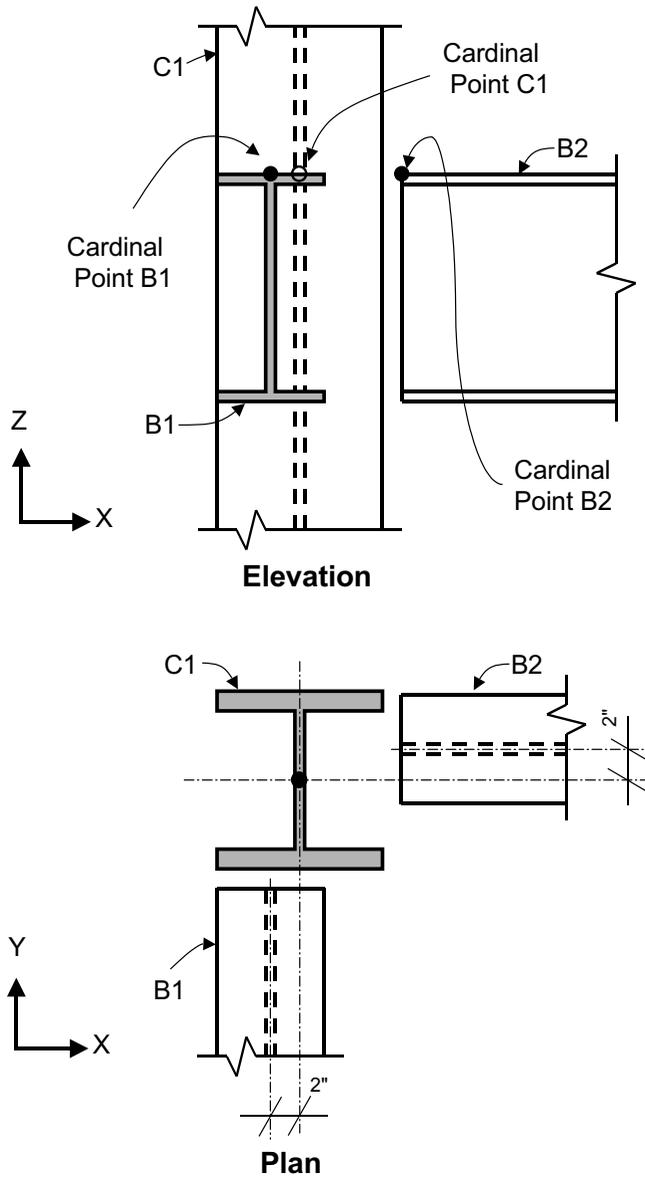


Figure 33
Example Showing Joint Offsets and Cardinal Points

End offsets are automatically calculated by the SAP2000 graphical interface for each element based on the maximum Section dimensions of all other elements that connect to that element at a common joint.

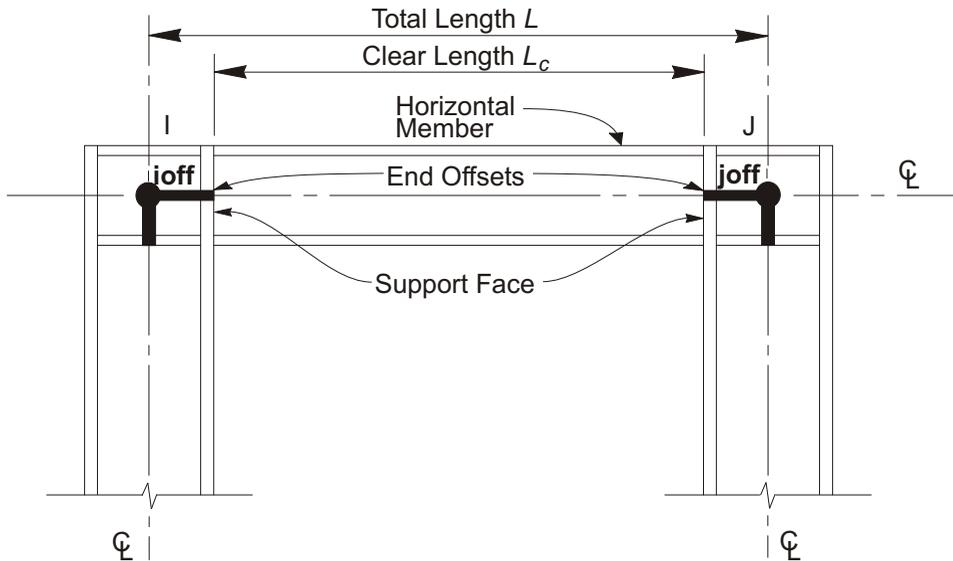


Figure 34
Frame Element End Offsets

Clear Length

The **clear length**, denoted L_c , is defined to be the length between the end offsets (support faces) as:

$$L_c = L - (\text{ioff} + \text{joff})$$

where L is the total element length. See Figure 34 (page 129).

If end offsets are specified such that the clear length is less than 1% of the total element length, the program will issue a warning and reduce the end offsets proportionately so that the clear length is equal to 1% of the total length. Normally the end offsets should be a much smaller proportion of the total length.

Rigid-end Factor

An analysis based upon the centerline-to-centerline (joint-to-joint) geometry of Frame elements may overestimate deflections in some structures. This is due to the stiffening effect caused by overlapping cross sections at a connection. It is more likely to be significant in concrete than in steel structures.

You may specify a **rigid-end factor** for each element using parameter **rigid**, which gives the fraction of each end offset that is assumed to be rigid for bending and shear deformation. The length **rigid ioff**, starting from joint **I**, is assumed to be rigid. Similarly, the length **rigid joff** is rigid at joint **j**. The flexible length L_f of the element is given by:

$$L_f = L \cdot \text{rigid}(\text{ioff}, \text{joff})$$

The rigid-zone offsets never affect axial and torsional deformation. The full element length is assumed to be flexible for these deformations.

The default value for **rigid** is zero. The maximum value of unity would indicate that the end offsets are fully rigid. You must use engineering judgment to select the appropriate value for this parameter. It will depend upon the geometry of the connection, and may be different for the different elements that frame into the connection. Typically the value for **rigid** would not exceed about 0.5.

Effect upon Non-prismatic Elements

At each end of a non-prismatic element, the Section properties are assumed to be constant within the length of the end offset. Section properties vary only along the clear length of the element between support faces. This is not affected by the value of the rigid-end factor, **rigid**.

See Subtopic “Non-prismatic Sections” (page 120) in this Chapter for more information.

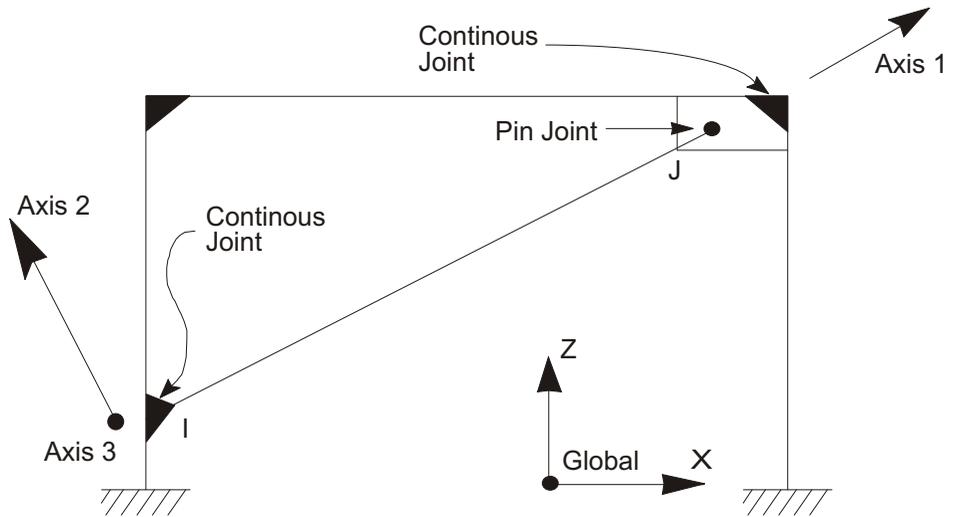
Effect upon Internal Force Output

All internal forces and moments are output at the faces of the supports and at other equally-spaced points within the clear length of the element. No output is produced within the end offset, which includes the joint. This is not affected by the value of the rigid-end factor, **rigid**.

See Topic “Internal Force Output” (page 144) in this Chapter for more information.

Effect upon End Releases

End releases are always assumed to be at the support faces, i.e., at the ends of the clear length of the element. If a moment or shear release is specified in either bending plane at either end of the element, the end offset is assumed to be rigid for bend-



For diagonal element: R3 is released at end J

Figure 35
Frame Element End Releases

ing and shear *in that plane at that end* (i.e., it acts as if **rigid** = 1). This does not affect the values of the rigid-end factor at the other end or in the other bending plane.

See Topic “End Releases” (page 131) in this Chapter for more information.

End Releases

Normally, the three translational and three rotational degrees of freedom at each end of the Frame element are continuous with those of the joint, and hence with those of all other elements connected to that joint. However, it is possible to release (disconnect) one or more of the element degrees of freedom from the joint when it is known that the corresponding element force or moment is zero. The releases are always specified in the element local coordinate system, and do not affect any other element connected to the joint.

In the example shown in Figure 35 (page 131), the diagonal element has a moment connection at End I and a pin connection at End J. The other two elements connecting to the joint at End J are continuous. Therefore, in order to model the pin condition the rotation R3 at End J of the diagonal element should be released. This assures that the moment is zero at the pin in the diagonal element.

Unstable End Releases

Any combination of end releases may be specified for a Frame element provided that the element remains stable; this assures that all load applied to the element is transferred to the rest of the structure. The following sets of releases are unstable, either alone or in combination, and are not permitted.

- Releasing U1 at both ends;
- Releasing U2 at both ends;
- Releasing U3 at both ends;
- Releasing R1 at both ends;
- Releasing R2 at both ends *and* U3 at either end;
- Releasing R3 at both ends *and* U2 at either end.

Effect of End Offsets

End releases are always applied at the support faces, i.e., at the ends of the element clear length. The presence of a moment or shear release will cause the end offset to be rigid in the corresponding bending plane at the corresponding end of the element.

See Topic “End Offsets” (page 127) in this Chapter for more information.

Named Property Sets

In addition to directly assigning end releases to frame elements, you can apply them to a frame element in a staged-construction Load Case using a Named Property Set of Frame Releases. A Named Property Set includes the same options that can be assigned to an element.

When property modifiers are changed in a staged construction Load Case, they do not change the response of the structure up to that stage, but only affect subsequent response. In other words, the effect is incremental. For example, consider a beam with assigned moment releases that is added between two columns in a staged construction case as follows:

- Stage 1: The beam is added to an existing structure and self-weight is applied. Because the element has moment releases assigned to it, the beam is added as simply-supported and has no fixed-end moments. The mid-span moment is 1000.

- Stage 2: Named Property Set “A” is applied that has no end releases. The mid-span moment does not change.
- Stage 3: Self-weight load is applied again (incrementally). The resulting mid-span moment increases to 1333, and the end span moments are -667 each.

This example illustrates the common case where beams are initially added as simply-supported, then connected to provide fully moment continuity. In another common situation, two independent staged construction cases can be considered: One for gravity load where certain members have moment releases, and the other for lateral load where the same members have moment continuity. Each of these cases can be used as the basis for further linear or nonlinear Load Cases, and the results considered together for design.

Nonlinear Properties

Two types of nonlinear properties are available for the Frame/Cable element: tension/compression limits and plastic hinges.

When nonlinear properties are present in the element, they only affect nonlinear analyses. Linear analyses starting from zero conditions (the unstressed state) behave as if the nonlinear properties were not present. Linear analyses using the stiffness from the end of a previous nonlinear analysis use the stiffness of the nonlinear property as it existed at the end of the nonlinear case.

Tension/Compression Limits

You may specify a maximum tension and/or a maximum compression that a frame/cable element may take. In the most common case, you can define a **no-compression** cable or brace by specifying the compression limit to be zero.

If you specify a tension limit, it must be zero or a positive value. If you specify a compression limit, it must be zero or a negative value. If you specify a tension and compression limit of zero, the element will carry no axial force.

The tension/compression limit behavior is elastic. Any axial extension beyond the tension limit and axial shortening beyond the compression limit will occur with zero axial stiffness. These deformations are recovered elastically at zero stiffness.

Bending, shear, and torsional behavior are not affected by the axial nonlinearity.

Plastic Hinge

You may insert plastic hinges at any number of locations along the clear length of the element. Detailed description of the behavior and use of plastic hinges is presented in Chapter “Frame Hinge Properties” (page 147).

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Frame element is lumped at the joints **I** and **J**. *No inertial effects are considered within the element itself.*

The total mass of the element is equal to the integral along the length of the mass density, **m**, multiplied by the cross-sectional area, **a**, plus the additional mass per unit length, **mpl**.

For non-prismatic elements, the mass varies linearly over each non-prismatic segment of the element, and is constant within the end offsets.

The total mass is apportioned to the two joints in the same way a similarly-distributed transverse load would cause reactions at the ends of a simply-supported beam. The effects of end releases are ignored when apportioning mass. The total mass is applied to each of the three translational degrees of freedom: UX, UY, and UZ. No mass moments of inertia are computed for the rotational degrees of freedom.

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties.”
- See Topic “Section Properties” (page 114) in this Chapter for the definition of **a** and **mpl**.
- See Subtopic “Non-prismatic Sections” (page 120) in this Chapter.
- See Topic “End Offsets” (page 127) in this Chapter.
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Static and Dynamic Analysis” (page 341).

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For a Frame element, the self-weight is a force that is distributed along the length of the

element. The magnitude of the self-weight is equal to the weight density, w , multiplied by the cross-sectional area, a , plus the additional weight per unit length, wpl .

For non-prismatic elements, the self-weight varies linearly over each non-prismatic segment of the element, and is constant within the end offsets.

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of w .
- See Topic “Section Properties” (page 114) in this Chapter for the definition of a and wpl .
- See Subtopic “Non-prismatic Sections” (page 120) in this Chapter.
- See Topic “End Offsets” (page 127) in this Chapter.
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Frame element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 134) in this Chapter for the definition of self-weight for the Frame element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Concentrated Span Load

The Concentrated Span Load is used to apply concentrated forces and moments at arbitrary locations on Frame elements. The direction of loading may be specified in

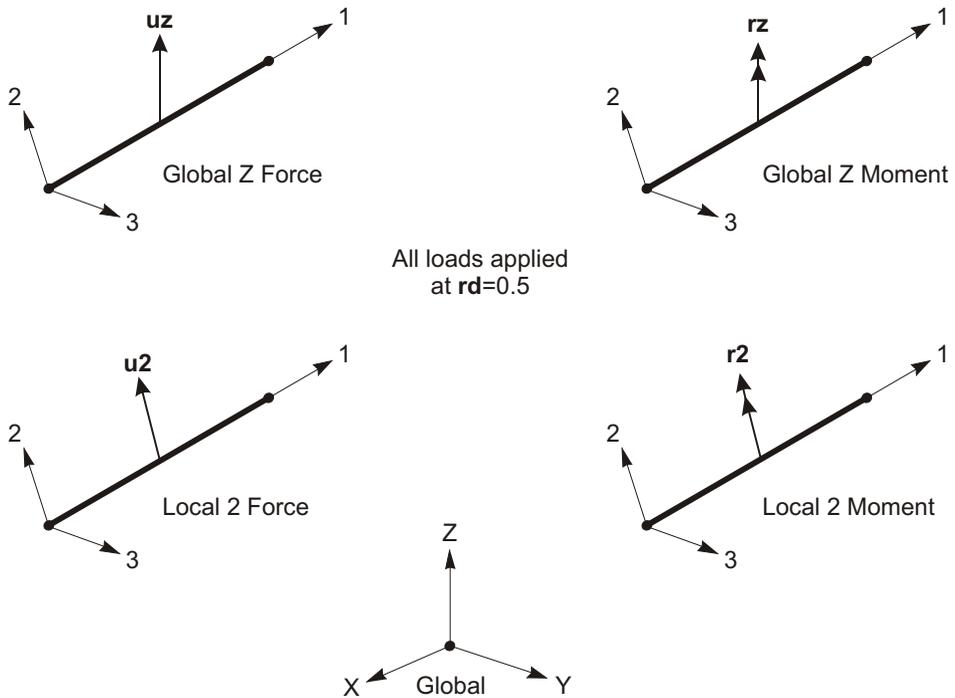


Figure 36
Examples of the Definition of Concentrated Span Loads

a fixed coordinate system (global or alternate coordinates) or in the element local coordinate system.

The location of the load may be specified in one of the following ways:

- Specifying a relative distance, **rd**, measured from joint **I**. This must satisfy $0 \leq \text{rd} \leq 1$. The relative distance is the fraction of element length;
- Specifying an absolute distance, **d**, measured from joint **I**. This must satisfy $0 \leq \text{d} \leq L$, where L is the element length.

Any number of concentrated loads may be applied to each element. Loads given in fixed coordinates are transformed to the element local coordinate system. See Figure 36 (page 136). Multiple loads that are applied at the same location are added together.

See Chapter “Load Patterns” (page 321) for more information.

Distributed Span Load

The Distributed Span Load is used to apply distributed forces and moments on Frame elements. The load intensity may be uniform or trapezoidal. The direction of loading may be specified in a fixed coordinate system (global or alternate coordinates) or in the element local coordinate system.

See Chapter “Load Patterns” (page 321) for more information.

Loaded Length

Loads may apply to full or partial element lengths. Multiple loads may be applied to a single element. The loaded lengths may overlap, in which case the applied loads are additive.

A loaded length may be specified in one of the following ways:

- Specifying two relative distances, **rda** and **rdb**, measured from joint **I**. They must satisfy $0 \leq \mathbf{rda} \leq \mathbf{rdb} \leq 1$. The relative distance is the fraction of element length;
- Specifying two absolute distances, **da** and **db**, measured from joint **I**. They must satisfy $0 \leq \mathbf{da} \leq \mathbf{db} \leq L$, where L is the element length;
- Specifying no distances, which indicates the full length of the element.

Load Intensity

The load intensity is a force or moment per unit of length. Except for the case of projected loads described below, the intensity is measured per unit of *element* length.

For each force or moment component to be applied, a single load value may be given if the load is uniformly distributed. Two load values are needed if the load intensity varies linearly over its range of application (a trapezoidal load).

See Figure 37 (page 138) and Figure 38 (page 139).

Projected Loads

A distributed snow or wind load produces a load intensity (force per unit of element length) that is proportional to the sine of the angle between the element and the direction of loading. This is equivalent to using a fixed load intensity that is measured

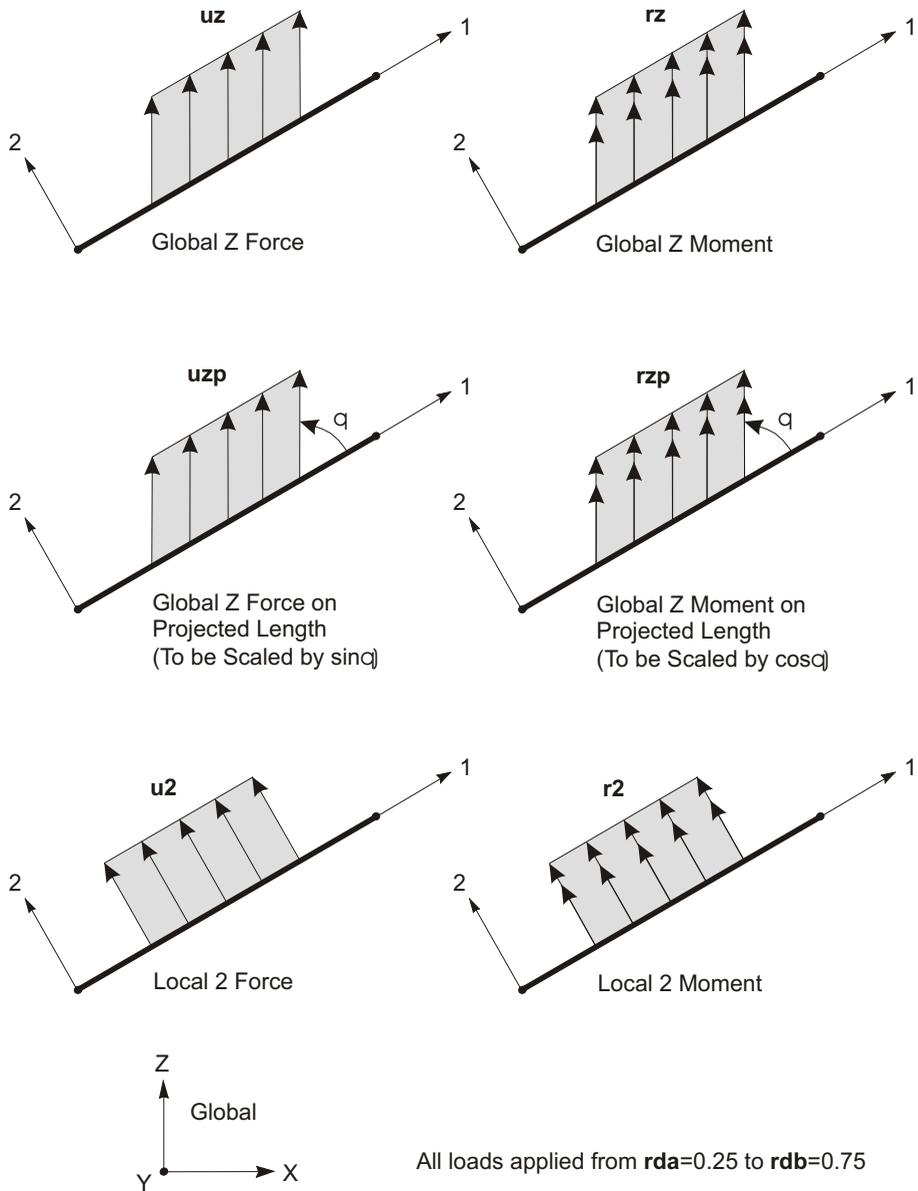


Figure 37
Examples of the Definition of Distributed Span Loads

per unit of *projected* element length. The fixed intensity would be based upon the

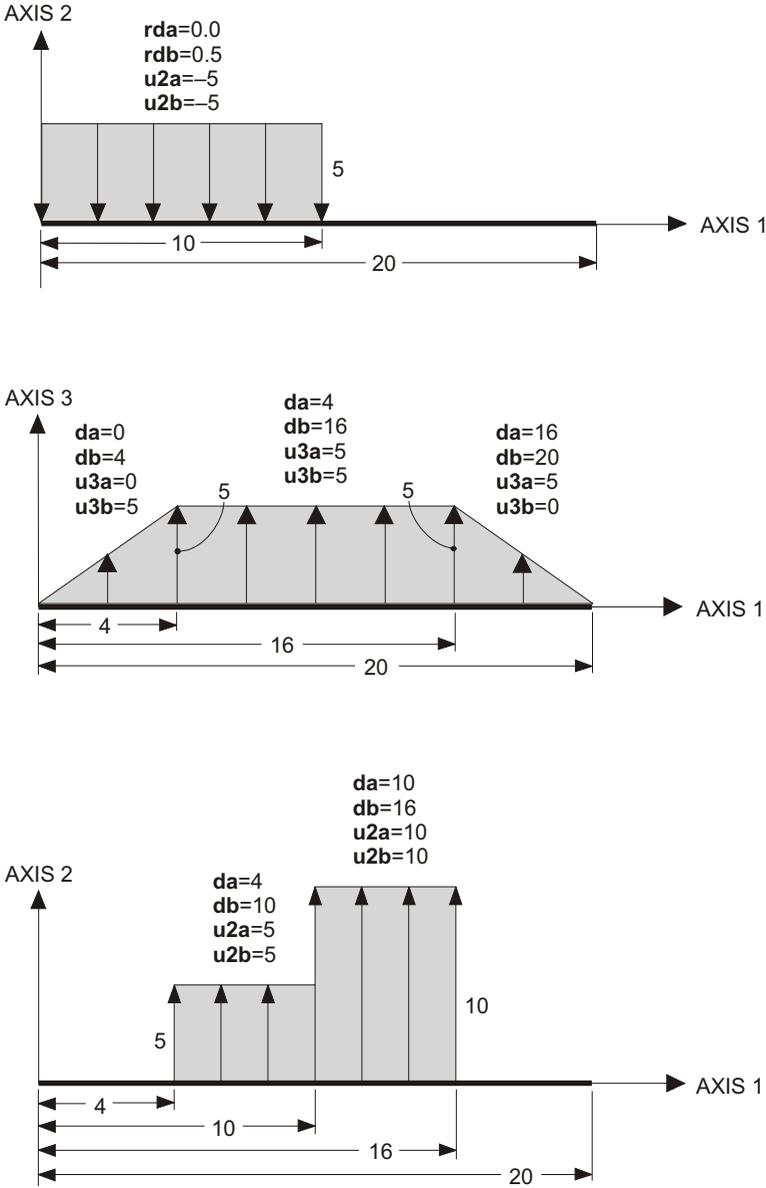


Figure 38
Examples of Distributed Span Loads

depth of snow or the wind speed; the projected element length is measured in a plane perpendicular to the direction of loading.

Distributed Span Loads may be specified as acting upon the projected length. The program handles this by reducing the load intensity according to the angle, θ , between the element local 1 axis and the direction of loading. Projected force loads are scaled by $\sin \theta$, and projected moment loads are scaled by $\cos \theta$. The reduced load intensities are then applied per unit of element length.

The scaling of the moment loads is based upon the assumption that the moment is caused by a force acting upon the projected element length. The resulting moment is always perpendicular to the force, thus accounting for the use of the cosine instead of the sine of the angle. The specified intensity of the moment should be computed as the product of the force intensity and the *perpendicular* distance from the element to the force. The appropriate sign of the moment must be given.

Temperature Load

Temperature Load creates thermal strain in the Frame element. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the element. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis.

Three independent Load Temperature fields may be specified:

- Temperature, **t**, which is constant over the cross section and produces axial strains
- Temperature gradient, **t2**, which is linear in the local 2 direction and produces bending strains in the 1-2 plane
- Temperature gradient, **t3**, which is linear in the local 3 direction and produces bending strains in the 1-3 plane

Temperature gradients are specified as the *change* in temperature per unit length. The temperature gradients are positive if the temperature increases (linearly) in the positive direction of the element local axis. The gradient temperatures are zero at the neutral axes, hence no axial strain is induced.

Each of the three Load Temperature fields may be constant along the element length or linearly interpolated from values given at the joints by a Joint Pattern.

See Chapter “Load Patterns” (page 321) for more information.

Strain Load

Six types of strain load are available, one corresponding to each of the internal forces and moments in a frame element. These are:

- Axial strain, ϵ , representing change in length per unit length. Positive strain increases the length of an unrestrained element, or causes compression in a restrained element.
- Shear strains, γ_2 and γ_3 , representing change in angle per unit length. The angle change is measured between the cross section and the neutral axis. Positive shear strain causes shear deformation in the same direction as would positive shear forces V2 and V3, respectively.
- Torsional curvature, θ_1 , representing change in torsional angle per unit length. Positive curvature causes deformation in the same direction as would positive torque T.
- Bending curvatures, θ_2 and θ_3 , representing change in angle per unit length. The angle is measured between adjacent sections that remain normal to the neutral axis. Positive curvature causes bending deformation in the same direction as would positive moments M2 and M3, respectively.

Each of the Strain Loads may be constant along the element length or linearly interpolated from values given at the joints by a Joint Pattern.

In an unrestrained element, strain loads cause deformation between the two ends of the element, but induce no internal forces. This unrestrained deformation has the same sign as would deformation caused by the corresponding (conjugate) forces and moments acting on the element. On the other hand, strain loading in a restrained element causes corresponding internal forces that have the opposite sign as the applied strain. Most elements in a real structure are connected to finite stiffness, and so strain loading would cause both deformation and internal forces. Note that the effects of shear and bending strain loading are coupled.

For more information, see Topic “Internal Force Output” (page 144) in this chapter, and also Chapter “Load Patterns” (page 321.)

Deformation Load

While Strain Load specifies a change in deformation per unit length, Deformation Load specifies the total deformation between the two ends of an unrestrained element. Deformation Load is internally converted to Strain Load, so you should

choose whichever type of loading is most conveniently specified for your particular application.

Currently only axial Deformation Load is available. The specified axial deformation is converted to axial Strain Load by simply dividing by the element length. The computed strain loads are assumed to be constant along the length of the element.

See Chapter “Load Patterns” (page 321) for more information.

Target-Force Load

Target-Force Load is a special type of loading where you specify a desired axial force, and deformation load is iteratively applied to achieve the target force. Since the axial force may vary along the length of the element, you must also specify the relative location where the desired force is to occur. Target-Force loading is only used for nonlinear static and staged-construction analysis. If applied in any other type of Load Case, it has no effect.

Unlike all other types of loading, target-force loading is not incremental. Rather, you are specifying the total force that you want to be present in the frame element at the end of the Load Case or construction stage. The applied deformation that is calculated to achieve that force may be positive, negative, or zero, depending on the force present in the element at the beginning of the analysis. When a scale factor is applied to a Load Pattern that contains Target-Force loads, the total target force is scaled. The increment of applied deformation that is required may change by a different scale factor.

See Topic “Target-Force Load” (page 331) in Chapter “Load Patterns” and Topic “Target-Force Iteration” (page 444) in Chapter “Nonlinear Static Analysis” for more information.

Internal Force Output

The **Frame element internal forces** are the forces and moments that result from integrating the stresses over an element cross section. These internal forces are:

- P, the axial force
- V2, the shear force in the 1-2 plane
- V3, the shear force in the 1-3 plane
- T, the axial torque

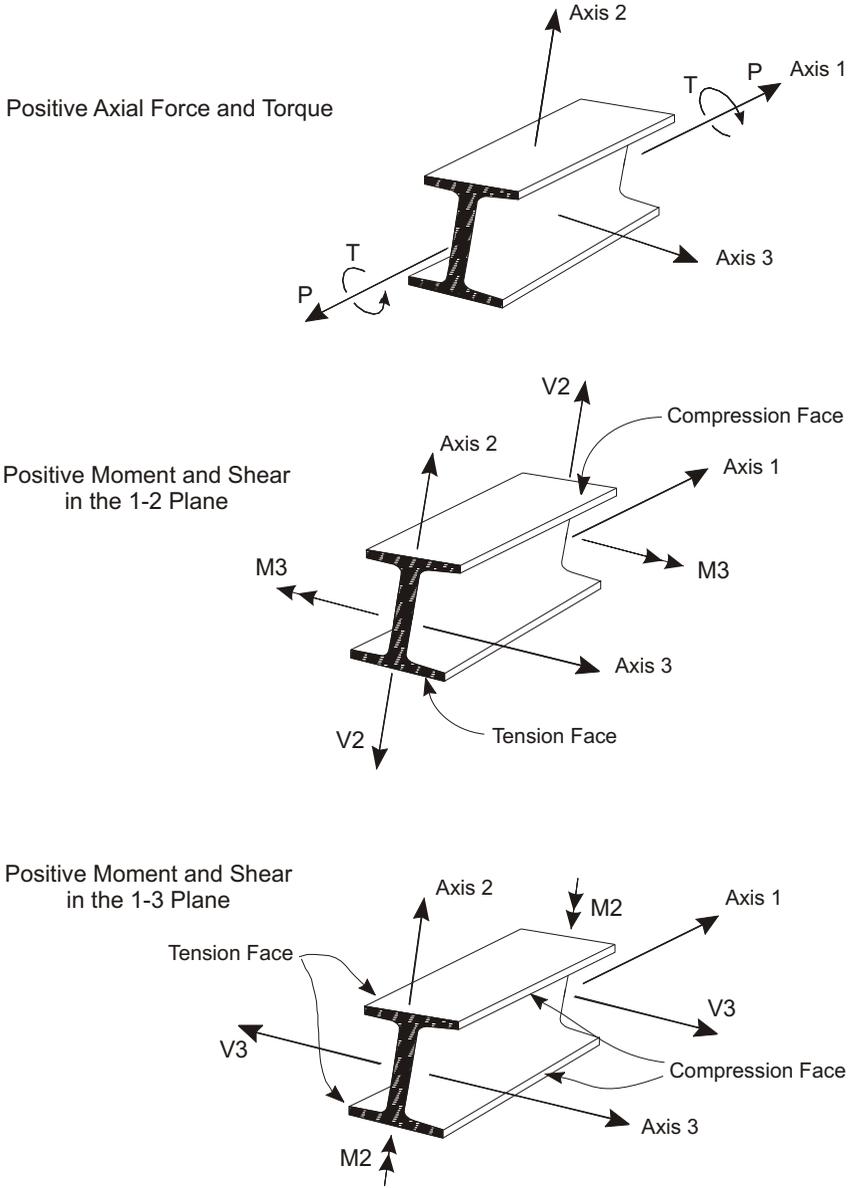


Figure 39
Frame Element Internal Forces and Moments

- M_2 , the bending moment in the 1-3 plane (about the 2 axis)
- M_3 , the bending moment in the 1-2 plane (about the 3 axis)

These internal forces and moments are present at every cross section along the length of the element, and may be requested as part of the analysis results.

The sign convention is illustrated in Figure 39 (page 143). Positive internal forces and axial torque acting on a positive 1 face are oriented in the positive direction of the element local coordinate axes. Positive internal forces and axial torque acting on a negative face are oriented in the negative direction of the element local coordinate axes. A positive 1 face is one whose outward normal (pointing away from element) is in the positive local 1 direction.

Positive bending moments cause compression at the positive 2 and 3 faces and tension at the negative 2 and 3 faces. The positive 2 and 3 faces are those faces in the positive local 2 and 3 directions, respectively, from the neutral axis.

Effect of End Offsets

When end offsets are present, internal forces and moments are output at the faces of the supports and at points within the clear length of the element. No output is produced within the length of the end offset, which includes the joint. Output will only be produced at joints **I** or **j** when the corresponding end offset is zero.

See Topic “End Offsets” (page 127) in this Chapter for more information.

Stress Output

Axial stress results are available for graphical display and tabular output along with the internal forces described above. The axial stress is denoted S11 and is computed at any point in the frame cross section as:

$$S11 = \frac{P}{a} + x2 \frac{M3}{i33} + x3 \frac{M2}{i22}$$

where

- P is the axial force, and M2 and M3 are the bending moments, as defined in Topic “Internal Force Output” (page 144)
- a is the cross-sectional area, and **i22** and **i33** are the section moments of inertia, as defined in Topic “Section Properties” (page 114)
- x2 and x3 are the coordinates of the point where the stress is calculated, measured from the centroid of the section along the local 2 and 3 axes, respectively.

Based on this definition, tensile stresses are always positive, and compressive stresses are always negative, regardless of the material.

Stresses are computed at the same axial stations as are the internal forces. At each station, stresses are computed at selected stress points over the cross section, with locations that depend upon the shape of the section:

- I-sections, T-sections, rectangles, boxes, channels, and angles at all corners where the maximum stresses could occur
- Cover-plated I sections at the same locations as for an I-section, plus the extreme corners of the cover plates, if present
- Circles and pipes at eight points on the circumference
- Section Designer sections at the stress points defined when the section was drawn; if no stress points have been defined, stress points are assumed at the four corners of the rectangular bounding box for the section; this box has dimensions $t_3 \times t_2$
- General sections and all other shapes at the four corners of the rectangular bounding box for the section; this box has dimensions $t_3 \times t_2$
- Nonprismatic sections computed as above from the interpolated shape, if the shape type is the same at both ends of the frame segment; if the shape type is not the same at both ends, then zero stress is reported.
- For all shapes except the box and pipe, stresses are also computed at the centroid of the section

In addition to the value of S11 at each stress point, two extreme stress values are reported at each station:

- S11Max – the maximum value over all stress points at that station
- S11Min – the minimum value over all stress points at that station

Stresses are computed for all load cases except for moving-load cases, for which zero values will be reported.

The following assumptions pertain to the stresses reported for frame elements:

- Stresses are computed for the base material of the section. No account is made for the modular ratio. This usually has no effect on stresses for any section types except for some Section Designer shapes.
- Stresses are computed based on the properties of the section assigned to the frame element. If the section property is changed during a staged-construction

load case, the calculated stress values may not be appropriate. However, the internal forces and moments are still correct.

- If any of the properties **a**, **i22**, or **i33** is zero, the stress *S11* will be reported as zero. It is recommended to use end releases rather than setting the section properties to zero.
- Property modifiers that are applied to **a**, **i22**, or **i33** do not change these property values when used for stress calculation. In some cases property modifiers may have an indirect effect upon the stresses if they affect the corresponding axial force or bending moments.

The description in this topic pertains only to the stress values reported as analysis results. Stresses used for frame design and bridge design are computed separately, as appropriate for the applicable material and design procedure.

Hinge Properties

You may insert plastic hinges at any number of locations along the clear length of any Frame element or Tendon object. ETABS also admits hinges in vertical shear-wall elements. Each hinge represents concentrated post-yield behavior in one or more degrees of freedom. Hinges only affect the behavior of the structure in nonlinear static and nonlinear time-history analyses.

Advanced Topics

- Overview
- Hinge Properties
- Automatic, User-Defined, and Generated Properties
- Automatic Hinge Properties
- Analysis Modeling
- Analysis Results

Overview

Yielding and post-yielding behavior can be modeled using discrete user-defined hinges. Hinges can be assigned to a frame element at any location along the clear

length of the element. Uncoupled moment, torsion, axial force and shear hinges are available. There are also coupled P-M2-M3 hinges which yield based on the interaction of axial force and bi-axial bending moments at the hinge location. Subsets of these hinges may include P-M2, P-M3, and M2-M3 behavior.

Fiber hinges P-M2-M3 can be defined, which are a collection of material points over the cross section. Each point represents a tributary area and has its own stress-strain curve. Plane sections are assumed to remain planar for the section, which ties together the behavior of the material points. Fiber hinges are often more realistic than force-moment hinges, but are more computationally intensive.

More than one type of frame hinge can exist at the same location, for example, you might assign both an M3 (moment) and a V2 (shear) hinge to the same end of a frame element. Hinge properties can be computed automatically from the element material and section properties according to FEMA-356 (FEMA, 2000) or ACSE 41-13 criteria.

For ETABS, hinges can also be assigned to vertical shear walls. These hinges are of type fiber P-M3, and always act at the center of the shell element. When hinges are present in a shear wall shell element, the vertical membrane stress behavior is governed by hinge, while horizontal and shear membrane stress, as well as out-of-plane bending behavior, are governed by the properties of the shell element.

Hinges only affect the behavior of the structure in nonlinear static and nonlinear time-history analyses. Hinge behavior does not affect nonlinear modal time-history (FNA) analyses unless the hinges are modeled as links, as described later in this chapter.

Strength loss is permitted in the hinge properties, and in fact the FEMA and ASCE hinges assume a sudden loss of strength. However, you should use this feature judiciously. Sudden strength loss is often unrealistic and can be very difficult to analyze, especially when elastic snap-back occurs. We encourage you to consider strength loss only when necessary, to use realistic negative slopes, and to carefully evaluate the results.

To help with convergence, the program automatically limits the negative slope of a hinge to be no stiffer than 10% of the elastic stiffness of the Frame element containing the hinge. If you need steeper slopes, you can assign a hinge overwrite that automatically meshes the frame element around the hinge. By reducing the size of the meshed element, you can increase the steepness of the drop-off.

Everything in this Chapter applies to Tendon objects as well as to Frame elements, although only the use of axial hinges makes sense for Tendons.

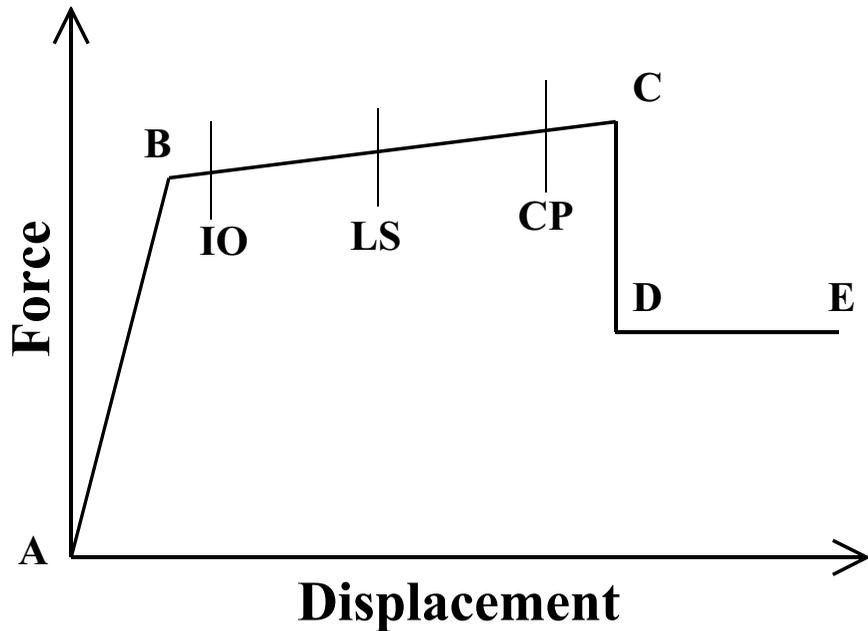


Figure 40

*The A-B-C-D-E curve for Force vs. Displacement
The same type of curve is used for Moment vs. Rotation*

Hinge Properties

A hinge property is a named set of nonlinear properties that can be assigned to points along the length of one or more Frame elements. You may define as many hinge properties as you need.

Force- and moment-type hinges are rigid-plastic. For each force degree of freedom (axial and shear), you may specify the plastic force-displacement behavior. For each moment degree of freedom (bending and torsion) you may specify the plastic moment-rotation behavior. Each hinge property may have plastic properties specified for any number of the six degrees of freedom. The axial force and the two bending moments may be coupled through an interaction surface. Degrees of freedom that are not specified remain elastic.

Fiber hinges are elastic-plastic and consist of a set of material points, each representing a portion of the frame cross-section having the same material. Force-deflection and moment-rotation curves are not specified, but rather are computed during the analysis from the stress-strain curves of the material points.

Hinge Length

Each plastic hinge is modeled as a discrete point hinge. All plastic deformation, whether it be displacement or rotation, occurs within the point hinge. This means you must assume a length for the hinge over which the plastic strain or plastic curvature is integrated.

There is no easy way to choose this length, although guidelines are given in FEMA-356 and ASCE 41-13. Typically it is a fraction of the element length, and is often on the order of the depth of the section, particularly for moment-rotation hinges.

You can approximate plasticity that is distributed over the length of the element by inserting many hinges. For example, you could insert ten hinges at relative locations within the element of 0.05, 0.15, 0.25, ..., 0.95, each with deformation properties based on an assumed hinge length of one-tenth the element length. Of course, adding more hinges will add more computational cost, so this should only be done where needed.

For force/moment-type hinges, elastic deformation occurs along the entire length of the Frame element and is not affected by the presence of the hinges. For fiber hinges, elastic behavior along the hinge length is determined from the hinge material stress-strain curves, and the elastic properties of the frame element are ignored within the hinge length. For this reason, the hinge length should not exceed the length of frame element.

Plastic Deformation Curve

For each force or moment degree of freedom, you define a force-displacement (moment-rotation) curve that gives the yield value and the plastic deformation following yield. This is done in terms of a curve with values at five points, A-B-C-D-E, as shown in Figure 40 (page 149). You may specify a symmetric curve, or one that differs in the positive and negative direction.

The shape of this curve as shown is intended for pushover analysis. You can use any shape you want. The following points should be noted:

- Point A is always the origin.
- Point B represents yielding. No deformation occurs in the hinge up to point B, regardless of the deformation value specified for point B. The displacement (rotation) at point B will be subtracted from the deformations at points C, D, and E. Only the plastic deformation beyond point B will be exhibited by the hinge.

- Point C represents the ultimate capacity for pushover analysis. However, you may specify a positive slope from C to D for other purposes.
- Point D represents a residual strength for pushover analysis. However, you may specify a positive slope from C to D or D to E for other purposes.
- Point E represent total failure. Beyond point E the hinge will drop load down to point F (not shown) directly below point E on the horizontal axis. If you do not want your hinge to fail this way, be sure to specify a large value for the deformation at point E.

You may specify additional deformation measures at points IO (immediate occupancy), LS (life safety), and CP (collapse prevention). These are informational measures that are reported in the analysis results and used for performance-based design. They do not have any effect on the behavior of the structure.

Prior to reaching point B, all deformation is linear and occurs in the Frame element itself, not the hinge. Plastic deformation beyond point B occurs in the hinge in addition to any elastic deformation that may occur in the element.

When the hinge unloads elastically, it does so without any plastic deformation, i.e., parallel to slope A-B.

Scaling the Curve

When defining the hinge force-deformation (or moment-rotation) curve, you may enter the force and deformation values directly, or you may enter normalized values and specify the scale factors that you used to normalized the curve.

In the most common case, the curve would be normalized by the yield force (moment) and yield displacement (rotation), so that the normalized values entered for point B would be (1,1). However, you can use any scale factors you want. They do not have to be yield values.

Remember that any deformation given from A to B is not used. This means that the scale factor on deformation is actually used to scale the plastic deformation from B to C, C to D, and D to E. However, it may still be convenient to use the yield deformation for scaling.

When automatic hinge properties are used, the program automatically uses the yield values for scaling. These values are calculated from the Frame section properties. See the next topic for more discussion of automatic hinge properties.

Strength Loss

Strength loss is permitted in the hinge properties, and in fact the FEMA hinges assume a sudden loss of strength. However, you should use this feature judiciously. Any loss of strength in one hinge causes load redistribution within the structure, possibly leading to failure of another hinge, and ultimately causing progressive collapse. This kind of analysis can be difficult and time consuming. Furthermore, any time negative stiffnesses are present in the model, the solution may not be mathematically unique, and so may be of questionable value.

Sudden strength loss (steep negative stiffness) is often unrealistic and can be even more difficult to analyze. When an unloading plastic hinge is part of a long beam or column, or is in series with any flexible elastic subsystem, “elastic snap-back” can occur. Here the elastic unloading deflection is larger than, and of opposite sign to, the plastic deformation. This results in the structure deflecting in the direction opposite the applied load. SAP2000, ETABS, and CSiBridge have a built-in mechanism to deal with snap-back for certain hinges, but this may not always be enough to handle several simultaneous snap-back hinge failures.

Consider carefully what you are trying to accomplish with your analysis. A well designed structure, whether new or retrofitted, should probably not have strength loss in any primary members. If an analysis shows strength loss in a primary member, you may want to modify the design and then re-analyze, rather than trying to push the analysis past the first failure of the primary members. Since you need to re-design anyway, what happens after the first failure is not relevant, since the behavior will become changed.

To help with convergence, the program automatically limits the negative slope of a hinge to be no stiffer than 10% of the elastic stiffness of the Frame element containing the hinge. By doing this, snap-back is prevented within the element, although it may still occur in the larger structure.

If you need steeper slopes, you can assign a Frame Hinge Overwrite that automatically meshes the Frame object around the hinge. When you assign this overwrite, you can specify what fraction of the Frame object length should be used for the element that contains the hinge. For example, consider a Frame object containing one hinge at each end, and one in the middle. If you assign a Frame Hinge Overwrite with a relative length of 0.1, the object will be meshed into five elements of relative lengths 0.05, 0.4, 0.1, 0.4, and 0.05. Each hinge is located at the center of an element with 0.1 relative length, but because two of the hinges fall at the ends of the object, half of their element lengths are not used. Because these elements are shorter than the object, their elastic stiffnesses are larger, and the program will permit larger negative stiffnesses in the hinges.

By reducing the size of the meshed element, you can increase the steepness of the drop-off, although the slope will never be steeper than you originally specified for the hinge. Again, we recommend gradual, realistic slopes whenever possible, unless you truly need to model brittle behavior.

Types of P-M2-M3 Hinges

Normally the hinge properties for each of the six degrees of freedom are uncoupled from each other. However, you have the option to specify coupled axial-force/bi-axial-moment behavior. This is called a P-M2-M3 or PMM hinge. Three types are available. In summary:

- Isotropic P-M2-M3 hinge: This hinge can handle complex and unsymmetrical PMM surfaces and can interpolate between multiple moment-rotation curves. Two-dimensional subsets of the hinge are available. It is limited to isotropic hysteresis, which may not be suitable for some structures.
- Parametric P-M2-M3 hinge: This hinge is limited to doubly symmetric section properties and uses a simple parametric definition of the PMM surface. Hysteretic energy degradation can be specified, making it more suitable than the isotropic hinge for extensive cyclic loading.
- Fiber P-M2-M3 hinge. This is the most realistic hinge, but may require the most computational resources in terms of analysis time and memory usage. Various hysteresis models are available and they can be different for each material in the hinge.

These hinges are described in more detail in the following topics.

Isotropic P-M2-M3 Hinge

This hinge can handle complex and unsymmetrical PMM surfaces and can interpolate between multiple moment-rotation curves. It is limited to isotropic hysteresis, which may not be suitable for some structures.

Three additional coupled hinges are available as subsets of the PMM hinge: P-M2, P-M3, and M2-M3 hinges.

Tension is Always Positive!

It is important to note that SAP2000 uses the sign convention where tension is always positive and compression is always negative, regardless of the material being used. This means that for some materials (e.g., concrete) the interaction surface may appear to be upside down.

Interaction (Yield) Surface

For the PMM hinge, you specify an interaction (yield) surface in three-dimensional P-M2-M3 space that represents where yielding first occurs for different combinations of axial force P, minor moment M2, and major moment M3.

The surface is specified as a set of P-M2-M3 curves, where P is the axial force (tension is positive), and M2 and M3 are the moments. For a given curve, these moments may have a fixed ratio, but this is not necessary. The following rules apply:

- All curves must have the same number of points.
- For each curve, the points are ordered from most negative (compressive) value of P to the most positive (tensile).
- The three values P, M2 and M3 for the first point of all curves must be identical, and the same is true for the last point of all curves
- When the M2-M3 plane is viewed from above (looking toward compression), the curves should be defined in a counter-clockwise direction
- The surface must be convex. This means that the plane tangent to the surface at any point must be wholly outside the surface. If you define a surface that is not convex, the program will automatically increase the radius of any points which are “pushed in” so that their tangent planes are outside the surface. A warning will be issued during analysis that this has been done.

You can explicitly define the interaction surface, or let the program calculate it using one of the following formulas:

- Steel, AISC-LRFD Equations H1-1a and H1-1b with $\phi = 1$
- Steel, FEMA-356 Equation 5-4
- Concrete, ACI 318-02 with $\phi = 1$

You may look at the hinge properties for the generated hinge to see the specific surface that was calculated by the program.

Moment-Rotation Curves

For PMM hinges you specify one or more moment/plastic-rotation curves corresponding to different values of P and moment angle θ . The moment angle is measured in the M2-M3 plane, where 0° is the positive M2 axis, and 90° is the positive M3 axis.

You may specify one or more axial loads P and one or more moment angles θ . For each pair (P, θ) , the moment-rotation curve should represent the results of the following experiment:

- Apply the fixed axial load P .
- Increase the moments M_2 and M_3 in a fixed ratio $(\cos \theta, \sin \theta)$ corresponding to the moment angle θ .
- Measure the plastic rotations R_{p2} and R_{p3} that occur after yield.
- Calculate the resultant moment $M = M_2 \cos \theta + M_3 \sin \theta$, and the projected plastic rotation $R_p = R_{p2} \cos \theta + R_{p3} \sin \theta$ at each measurement increment
- Plot M vs. R_p , and supply this data to SAP2000

Note that the measured direction of plastic strain may not be the same as the direction of moment, but the projected value is taken along the direction of the moment. In addition, there may be measured axial plastic strain that is not part of the projection. However, during analysis the program will recalculate the total plastic strain based on the direction of the normal to the interaction (yield) surface.

During analysis, once the hinge yields for the first time, i.e., once the values of P , M_2 and M_3 first reach the interaction surface, a net moment-rotation curve is interpolated to the yield point from the given curves. This curve is used for the rest of the analysis for that hinge.

If the values of P , M_2 , and M_3 change from the values used to interpolate the curve, the curve is adjusted to provide an energy equivalent moment-rotation curve. This means that the area under the moment-rotation curve is held fixed, so that if the resultant moment is smaller, the ductility is larger. This is consistent with the underlying stress strain curves of axial “fibers” in the cross section.

As plastic deformation occurs, the yield surface changes size according to the shape of the M - R_p curve, depending upon the amount of plastic work that is done. You have the option to specify whether the surface should change in size equally in the P , M_2 , and M_3 directions, or only in the M_2 and M_3 directions. In the latter case, axial deformation behaves as if it is perfectly plastic with no hardening or collapse. Axial collapse may be more realistic in some hinges, but it is computationally difficult and may require nonlinear direct-integration time-history analysis if the structure is not stable enough to redistribute any dropped gravity load.

Parametric P-M2-M3 Hinge

This hinge is limited to doubly symmetric section properties and uses a simple parametric definition of the PMM surface. Hysteretic energy degradation can be specified, making it more suitable than the isotropic hinge for extensive cyclic loading.

Two versions of the hinge are available, one for steel frame sections, and one for re-inforced-concrete frame sections. Currently this hinge is only available in ETABS, and will be added to SAP2000 and CSiBridge in subsequent versions.

The description and theory for this hinge formulation are presented in the Technical Note “Parametric P-M2-M3 Hinge Model”. This document can be found in the Manuals subfolder where the software is installed on your computer. It can be accessed from inside the software using the menu command Help > Documentation > Technical Notes.

Detailed descriptions of the input values needed to define the properties for either the steel or concrete hinge are available from the Help facility within the software. This can be accessed using the menu command Help > Product Help, or pressing the F1 key at any time.

Fiber P-M2-M3 Hinge

The Fiber P-M2-M3 (Fiber PMM) hinge models the axial behavior of a number of representative axial “fibers” distributed across the cross section of the frame element. Each fiber has a location, a tributary area, and a stress-strain curve. The axial stresses are integrated over the section to compute the values of P, M2 and M3. Likewise, the axial deformation U1 and the rotations R2 and R3 are used to compute the axial strains in each fiber. Plane sections are assumed to remain planar.

You can define your own fiber hinge, explicitly specifying the location, area, material and its stress-strain curve for each fiber, or you can let the program automatically create fiber hinges for circular and rectangular frame sections.

The Fiber PMM hinge is more “natural” than the Isotropic or Parametric PMM hinges described above, since it automatically accounts for interaction, changing moment-rotation curve, and plastic axial strain. However, it is also more computationally intensive, requiring more computer storage and execution time. You may have to experiment with the number of fibers needed to get an optimum balance between accuracy and computational efficiency.

Strength loss in a fiber hinge is determined by the strength loss in the underlying stress-strain curves. Because all the fibers in a cross section do not usually fail at the same time, the overall hinges tend to exhibit more gradual strength loss than hinges with directly specified moment-rotation curves. This is especially true if reasonable hinge lengths are used. For this reason, the program does not automatically restrict the negative drop-off slopes of fiber hinges. However, we still recommend that you pay close attention to the modeling of strength loss, and modify the stress-strain curves if necessary.

For more information:

- See Topic “Stress-Strain Curves” (page 80) in Chapter “Material Properties.”
- See Topic “Section-Designer Sections” (page 118) Chapter “The Frame Element.”

Hysteresis Models

The plastic force-deformation or moment-rotation curve defines the nonlinear behavior under monotonic loading. This curve, combined with the elastic behavior of the hinge length in the parent frame element, is also known as the backbone curve for the hinge.

Under load reversal or cyclic loading, the behavior will deviate from the backbone curve. Several different hysteresis models are available to describe this behavior for different types of materials. For the most part, these differ in the amount of energy they dissipate in a given cycle of deformation, and how the energy dissipation behavior changes with an increasing amount of deformation.

Hysteresis models are described in Topic “Hysteresis Models” (page 85) of Chapter “Material Properties.”

Hysteresis models are applicable to the different types of hinges as follows:

- Single degree of freedom hinges: All inelastic models (kinematic, degrading, Takeda, pivot, concrete, BRB hardening, and isotropic)
- Coupled P-M2-M3, P-M2, P-M3, and M2-M3 hinges: Isotropic model only
- Fiber P-M2-M3 hinges: For each material fiber, all models (elastic, kinematic, degrading, Takeda, pivot, concrete, BRB hardening, and isotropic)

Note that all of these models are available in the current version of ETABS. Some of the models are not yet available in SAP2000 and CSiBridge but will be added in subsequent versions.

Hysteretic behavior may affect nonlinear static and nonlinear time-history load cases that exhibit load reversals and cyclic loading. Monotonic loading is not affected. Note, however, that even static pushover load cases can produce load reversal in some hinges caused by strength loss in other hinges.

Automatic, User-Defined, and Generated Properties

There are three types of hinge properties in SAP2000:

- Automatic hinge properties
- User-defined hinge properties
- Generated hinge properties

Only automatic hinge properties and user-defined hinge properties can be assigned to frame elements. When automatic or user-defined hinge properties are assigned to a frame element, the program automatically creates a generated hinge property for each and every hinge.

The built-in automatic hinge properties for steel members are based on Table 5-6 in FEMA-356. The built-in automatic hinge properties for concrete members are based on Tables 6-7 and 6-8 in FEMA-356, or on Caltrans specifications for concrete columns. After assigning automatic hinge properties to a frame element, the program generates a hinge property that includes specific information from the frame section geometry, the material, and the length of the element. You should review the generated properties for their applicability to your specific project.

User-defined hinge properties can either be based on a hinge property generated from automatic property, or they can be fully user-defined.

A generated property can be converted to user-defined, and then modified and re-assigned to one or more frame elements. This way you can let the program do much of the work for you using automatic properties, but you can still customize the hinges to suit your needs. However, once you convert a generated hinge to user-defined, it will no longer change if you modify the element, its section or material.

It is the generated hinge properties that are actually used in the analysis. They can be viewed, but they can not be modified. Generated hinge properties have an automatic naming convention of LabelH#, where Label is the frame element label, H stands for hinge, and # represents the hinge number. The program starts with hinge number 1 and increments the hinge number by one for each consecutive hinge applied to the frame element. For example if a frame element label is F23, the gener-

ated hinge property name for the second hinge assigned to the frame element is F23H2.

The main reason for the differentiation between defined properties (in this context, defined means both automatic and user-defined) and generated properties is that typically the hinge properties are section dependent. Thus it would be necessary to define a different set of hinge properties for each different frame section type in the model. This could potentially mean that you would need to define a very large number of hinge properties. To simplify this process, the concept of automatic properties is used in SAP2000. When automatic properties are used, the program combines its built-in default criteria with the defined section properties for each element to generate the final hinge properties. The net effect of this is that you do significantly less work defining the hinge properties because you don't have to define each and every hinge.

Automatic Hinge Properties

Automatic hinge properties are based upon a simplified set of assumptions that may not be appropriate for all structures. You may want to use automatic properties as a starting point, and then convert the corresponding generated hinges to user-defined and explicitly overwrite calculated values as needed.

Automatic properties require that the program have detailed knowledge of the Frame Section property used by the element that contains the hinge. For this reason, only the following types of automatic hinges are available:

Concrete Beams in Flexure

M2 or M3 hinges can be generated using FEMA Table 6-7 (I) for the following shapes:

- Rectangle
- Tee
- Angle
- Section Designer

Concrete Columns in Flexure

M2, M3, M2-M3, P-M2, P-M3, or P-M2-M3 hinges can be generated using FEMA Table 6-8 (I), for the following shapes:

- Rectangle
- Circle
- Section Designer

or using Caltrans specifications, for the following shapes:

- Section Designer only

Steel Beams in Flexure

M2 or M3 hinges can be generated using FEMA Table 5-6, for the following shapes:

- I/Wide-flange only

Steel Columns in Flexure

M2, M3, M2-M3, P-M2, P-M3, or P-M2-M3 hinges can be generated using FEMA Table 5-6, for the following shapes:

- I/Wide-flange
- Box

Steel Braces in Tension/Compression

P (axial) hinges can be generated using FEMA Table 5-6, for the following shapes:

- I/Wide-flange
- Box
- Pipe
- Double channel
- Double angle

Fiber Hinge

P-M2-M3 hinges can be generated for steel or reinforced concrete members using the underlying stress-strain behavior of the material for the following shapes:

- Rectangle
- Circle

Additional Considerations

You must make sure that all required design information is available to the Frame section as follows:

- For concrete Sections, the reinforcing steel must be explicitly defined, or else the section must have already been designed by the program before nonlinear analysis is performed
- For steel Sections, Auto-select Sections can only be used if they have already been designed so that a specific section has been chosen before nonlinear analysis is performed

For more information, see the on-line help that is available while assigning automatic hinges to Frame elements in the Graphical User Interface.

Analysis Modeling

Hinges are assigned to a Frame or Shell (shear wall) element to represent the nonlinear behavior of their parent element. When the analysis model is created, there are two ways the hinge can be represented:

- Hinge embedded in the element
- Hinge as a separate link element

The latter method is currently only available in the ETABS Ultimate level, and enables hinge behavior to be considered in nonlinear model time-history (FNA) load cases. As a rule, FNA analysis runs significantly faster than nonlinear direct-integration time-history analysis. Nonlinear static analysis and nonlinear direct-integration time history analysis are available for both types of analysis modeling.

When the hinge is modeled as a link element, the parent Frame element is divided at the hinge location into separate subelements, and a zero-length link element is created that contains the hinge property and connects the frame subelements. A very small amount of axial mass and rotational inertia are added at the two connecting joints to improve FNA iteration. A similar internal modeling is employed for shear wall elements when the hinge is modeled as a link element.

A second, independent modeling option is available to assign automatic subdivision of Frame elements at hinge locations. Using this assignment, you specify a relative length that is used when creating the analysis model of the hinges for the selected elements. The effect of this depends upon how the hinge is modeled:

- For the hinge embedded in the element: The Frame object is subdivided into separate frame elements, with one element containing the hinge that is equal in length to that specified in the assignment.

This has the advantage of introducing more degrees of freedom into the model that may improve convergence when multiple hinges are failing at the same time, with a possible increase in computation time. In addition, steeper drop-offs are permitted when the hinge curve exhibits strength loss because the element containing the hinge is shorter, and hence stiffer.

On the other hand, not subdividing the frame element leads to a smaller analysis model, typically requiring less computation time and storage. In addition, stiffness proportional damping for nonlinear direct-integration time-history analysis is better modeled in longer elements.

- For the hinge as a separate link element: The subdivision into two frame elements and a zero-length link is not changed. However, the elastic flexibility of the link is changed to be equal to the length of the frame element specified in the assignment, and the corresponding length of the adjacent frame sub-elements are made rigid.

This has the advantage of improving stiffness-proportional damping in nonlinear direct-integration time-history analysis, and can be recommended for this reason. On the other hand, this is not necessary for FNA analysis.

The default relative length for automatic subdivision is 0.02. Recommended values typically range from 0.02 to 0.25.

Computational Considerations

The most important advice is to only add hinges to the model where nonlinear behavior is expected to have a significant effect on the analysis and design. Adding extra hinges increases the time and effort it takes to create the model, to run the analyses, and to interpret the results.

Start with the simplest model possible so that you can make many analysis runs quickly. This helps to better understand the behavior of your structure early in the design process and to correct modeling errors. Add hinges and complexity gradually as you determine where nonlinearity is expected and/or desired.

Adding hinges everywhere to find the nonlinearity is tempting, but this approach usually wastes much more time than incrementally growing the model.

Most models with hinges benefit from using event-to-event stepping for nonlinear static and nonlinear direct-integration time-history load cases. This is particularly true for the parametric P-M2-M3 hinge. However, it may be necessary to turn off event-to-event stepping if the model has a very large number of hinges, or if there is a significant amount of other types of nonlinearity in the structure. This is best determined by running analyses both with and without events to see which is most efficient.

Most nonlinear time-history analysis benefits from the presence of mass at the nonlinear degrees-of-freedom. Inertia tends to stabilize iteration when the nonlinear behavior is changing rapidly. This is particularly true for FNA analysis. For ETABS, it is usually best to define the mass source to include vertical mass and to *not* lump the mass at the story levels for models that have hinges.

For FNA analysis, it is usually most efficient to damp out the very high modes. Some of the Ritz modes needed for FNA analysis can be expected to be of high frequency. An example of how to do this would be to define the load-case damping to be of type “Interpolated by Frequency”. Then specify your desired structure damping ratio (say 0.025) for frequencies up to 999 Hz, and a damping ratio of 0.99 for frequencies above 1000 Hz. You can experiment with this cutoff value to see the effect on runtime and results.

Analysis Results

For each output step in a nonlinear static or nonlinear direct-integration time-history Load Case, you may request analysis results for the hinges. These results include:

- The forces and/or moments carried by the hinge. Degrees of freedom not defined for the hinge will report zero values, even though non-zero values are carried rigidly through the hinge.
- The plastic displacements and/or rotations.
- The most extreme state experienced by the hinge in any degree of freedom. This state does not indicate whether it occurred for positive or negative deformation:
 - A to B
 - B to C
 - C to D
 - D to E

- > E
- The most extreme performance status experienced by the hinge in any degree of freedom. This status does not indicate whether it occurred for positive or negative deformation:
 - A to B
 - B to IO
 - IO to LS
 - LS to CP
 - > CP

When you display the deflected shape in the graphical user interface for a nonlinear static or nonlinear direct-integration time-history Load Case, the hinges are plotted as colored dots indicating their most extreme state or status:

- B to IO
- IO to LS
- LS to CP
- CP to C
- C to D
- D to E
- > E

The colors used for the different states are indicated on the plot. Hinges that have not experienced any plastic deformation (A to B) are not shown.

Chapter IX

The Cable Element

The **Cable** element is a highly nonlinear element used to model the catenary behavior of slender cables under their own self-weight. Tension-stiffening and large-deflections nonlinearity are inherently included in the formulation. Nonlinear analysis is required to make use of the Cable element. Linear analyses can be performed that use the stiffness from the end of nonlinear Load Cases.

Advanced Topics

- Overview
- Joint Connectivity
- Undeformed Length
- Shape Calculator
- Degrees of Freedom
- Local Coordinate System
- Section Properties
- Property Modifiers
- Mass
- Self-Weight Load

- Gravity Load
- Distributed Span Load
- Temperature Load
- Strain and Deformation Load
- Target-Force Load
- Nonlinear Analysis
- Element Output

Overview

The Cable element uses an elastic catenary formulation to represent the behavior of a slender cable under its own self-weight, temperature, and strain loading. This behavior is highly nonlinear, and inherently includes tension-stiffening (P-delta) and large-deflection effects. Slack and taut behavior is automatically considered.

In the graphical user interface, you can draw a cable object connecting any two points. A shape calculator is available to help you determine the undeformed length of the cable. The undeformed length is extremely critical in determining the behavior of the cable.

An unloaded, slack cable is not stable and has no unique position. Therefore linear Load Cases that start from zero initial conditions may be meaningless. Instead, all linear Load Cases should use the stiffness from the end of a nonlinear static Load Case in which all cables are loaded by their self-weight or other transverse load. For cases where no transverse load is present on a slack Cable element, the program will internally assume a very small self-weight load in order to obtain a unique shape. However, it is better if you apply a realistic load for this purpose.

Each Cable element may be loaded by gravity (in any direction), distributed forces, strain and deformation loads, and loads due to temperature change. To apply concentrated loads, a cable should be divided at the point of loading, and the force applied to the connecting joint.

Target-force loading is available that iteratively applies deformation load to the cable to achieve a desired tension.

Element output includes the axial force and deflected shape at a user-specified number of equally-spaced output stations along the length of the element.

You have the option when drawing a cable object in the model to use the catenary element of this chapter, or to model the cable as a series of straight frame elements. Using frame elements allows you to consider material nonlinearity and complicated loading, but the catenary formulation is better suited to most applications.

Joint Connectivity

A Cable element is represented by a curve connecting two joints, **I** and **j**. The two joints must not share the same location in space. The two ends of the element are denoted end I and end J, respectively.

The shape of the cable is defined by undeformed length of the cable and the load acting on it, unless it is taut with no transverse load, in which case it is a straight line.

Undeformed Length

In the graphical user interface, you can draw a cable object connecting any two points. A shape calculator is available to help you determine the undeformed length of the cable. The relationship between the undeformed length and the chord length (the distance between the two end joints) is extremely critical in determining the behavior of the cable.

In simple terms, when the undeformed length is longer than the chord length, the cable is slack and has significant sag. When the undeformed length is shorter than the chord length, the cable is taut and carries significant tension with little sag.

When transverse load acts on the cable, there is a transition range where the undeformed length is close to the chord length. In this regime, the tension and sag interact in a highly nonlinear way with the transverse load.

Temperature, strain, and distortion loads can change the length of the cable. The effect of these changes is similar to changing the undeformed length, except that they do not change the weight of the cable. Strain in the cable due to any source is calculated as the difference between the total length and the undeformed length, divided by the undeformed length (engineering strain).

If the undeformed length of a cable is shorter than the chord length at the beginning of a nonlinear analysis, or when the cable is added to the structure during staged construction, tension will immediately exist in the cable and iteration may be required to bring the structure into equilibrium before any load is applied.

Shape Calculator

The ultimate purpose of the shape calculator (also called Cable Layout form) in the graphical user interface is to help you calculate the undeformed length of a cable object. By default, the undeformed length is assumed to be equal to the chord length between the undeformed positions of the two end joints.

You may specify a vertical load acting on the cable consisting of:

- Self-weight (always included in the shape calculator)
- Additional weight per unit of undeformed length of the cable
- Addition load per unit horizontal length between the two joints

Note that these loads are *only* used in the shape calculator. They are *not* applied to the element during analysis. Loads to be used for analysis must be assigned to the elements in Load Patterns.

You may choose one of the following ways to calculate the undeformed length:

- Specifying the undeformed length, either absolute or relative to the chord length
- Specifying the maximum vertical sag, measured from the chord to the cable
- Specifying the maximum low-point sag, measured from the joint with the lowest Z elevation to the lowest point on the cable
- Specifying the constant horizontal component of tension in the cable
- Specifying the tension at either end of the cable
- Requesting the shape which gives the minimum tension at either end of the cable

See Figure 41 (page 169) for a description of the cable geometry.

Note that there does exist an undeformed length that yields a minimum tension at either end of the cable. Longer cables carry more self weight, increasing the tension. Shorter cables are tauter, also increasing the tension. If you intend to specify the tension at either end, it is a good idea first to determine what is the minimum tension, since attempts to specify a lower tension will fail. When a larger value of tension is specified, the shorter solution will be returned.

It is important to note that the shape calculated here may not actually occur during any Load Case, nor are the tensions calculated here directly imposed upon the cable. Only the cable length is determined. The deformed shape of the cable and the

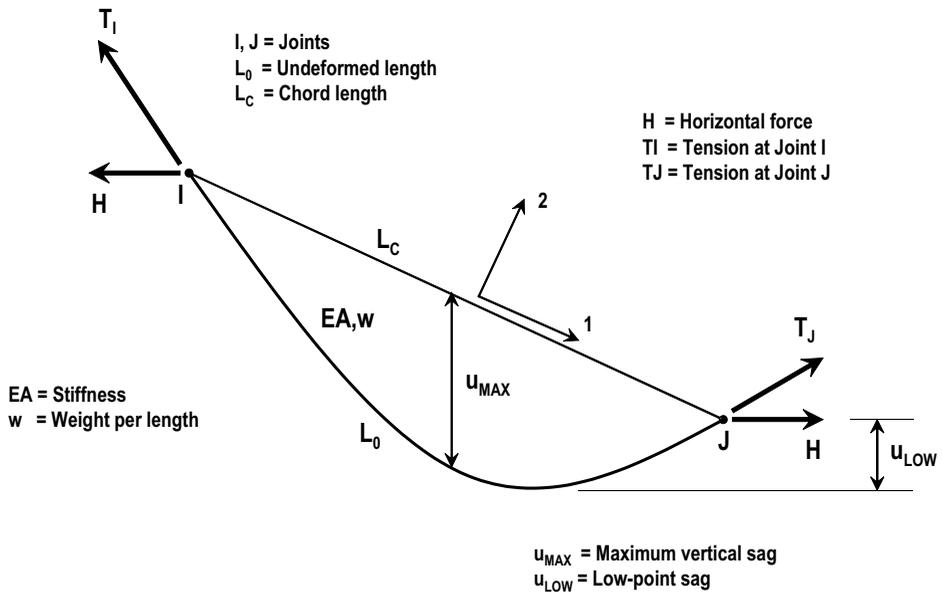


Figure 41

Cable Element, showing connectivity, local axes, dimensions, properties, and shape parameters

tensions it carries will depend upon the loads applied and the behavior of the structure during analysis. For example, the shape calculator assumes that the two end joints remain fixed. However, if the cable is connected to a deforming structure, the chord length and its orientation may change, yielding a different solution.

Cable vs. Frame Elements

In the shape calculator, you may specify whether the cable is to be modeled with the catenary element of this chapter, or using straight frame elements.

If you are interested in highly variable loading or material nonlinearity, using frame elements may be appropriate. Large-deflection geometrically nonlinear analysis of the entire structure will be needed to capture full cable behavior. P-delta analysis

with compression limits may be sufficient for some applications. For more information, see Chapter “The Frame Element” (page 105).

For most cable applications, however, the catenary cable element is a better choice, especially if the cable is very slender, or significant support movement is expected. Nonlinear analysis is still required, but the geometric nonlinearity (P-delta and/or large-deflection behavior) of the catenary element will be considered internally regardless of how the rest of the structure is treated.

Number of Segments

In the shape calculator, you may specify the number of segments into which the cable object should be broken. Each segment will be modeled as a single catenary cable or single frame element.

For the catenary element, a single segment is usually the best choice unless you are considering concentrated loads or intermediate masses for cable vibration.

For the frame element, multiple segments (usually at least eight, and sometimes many more) are required to capture the shape variation, unless you are modeling a straight stay or brace, in which case a single segment may suffice.

For more information, see Chapter “Objects and Elements” (page 7)

Degrees of Freedom

The Cable element activates the three translational degrees of freedom at each of its connected joints. Rotational degrees of freedom are not activated. This element contributes stiffness to all of these translational degrees of freedom.

For more information, see Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom.”

Local Coordinate System

Each Cable element has its own **element local coordinate system** which can be used to define loads acting on the element. The axes of this local system are denoted 1, 2 and 3. The first axis is directed along the chord connecting the two joints of the element; the remaining two axes lie in the plane perpendicular to the chord with an orientation that you specify. This coordinate system *does not* necessarily corre-

spond to the direction of sag of the cable, and *does not* change as the direction of sag changes during loading.

The definition of the cable element local coordinate system is not usually important unless you want to apply concentrated or distributed span loads in the element local system.

The definition of the Cable local coordinate system is exactly the same as for the Frame element. For more information, see Topics “Local Coordinate System” (page 108) and “Advanced Local Coordinate System” (page 110) in Chapter “The Frame Element.”

Section Properties

A **Cable Section** is a set of material and geometric properties that describe the cross-section of one or more Cable elements. Sections are defined independently of the Cable elements, and are assigned to the elements.

Cable Sections are always assumed to be circular. You may specify either the diameter or the cross-sectional area, from which the other value is computed. Bending moments of inertia, the torsional constant, and shear areas are also computed by the program for a circular shape.

Material Properties

The material properties for the Section are specified by reference to a previously-defined Material. Isotropic material properties are used, even if the Material selected was defined as orthotropic or anisotropic. The material properties used by the Section are:

- The modulus of elasticity, **e1**, for axial stiffness
- The coefficient of thermal expansion, **a1**, for temperature loading
- The mass density, **m**, for computing element mass
- The weight density, **w**, for computing Self-Weight and Gravity Loads

The material properties **e1** and **a1** are obtained at the material temperature of each individual Cable element, and hence may not be unique for a given Section. See Chapter “Material Properties” (page 69) for more information.

Geometric Properties and Section Stiffnesses

For the catenary formulation, the section has only axial stiffness, given by $\mathbf{a} \mathbf{e1}$, where \mathbf{a} is the cross-sectional area and $\mathbf{e1}$ is the modulus of elasticity.

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Cable element is lumped at the joints **I** and **J**. *No inertial effects are considered within the element itself.*

The total mass of the element is equal to the undeformed length of the element multiplied by the mass density, \mathbf{m} , and by the cross-sectional area, \mathbf{a} . It is apportioned equally to the two joints. The mass is applied to each of the three translational degrees of freedom: UX, UY, and UZ.

To capture dynamics of a cable itself, it is necessary to divide the cable object into multiple segments. A minimum of four segments is recommended for this purpose. For many structures, cable vibration is not important, and no subdivision is necessary.

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties.”
- See Topic “Section Properties” (page 171) in this Chapter for the definition of \mathbf{a} .
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Static and Dynamic Analysis” (page 341).

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For a Cable element, the self-weight is a force that is distributed along the arc length of the element. The magnitude of the self-weight is equal to the weight density, \mathbf{w} , multiplied by the cross-sectional area, \mathbf{a} . As the cable stretches, the magnitude is correspondingly reduced, so that the total load does not change.

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of w .
- See Topic “Section Properties” (page 171) in this Chapter for the definition of a .
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Cable element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element. The magnitude of a unit gravity load is equal to the weight density, w , multiplied by the cross-sectional area, a . As the cable stretches, the magnitude is correspondingly reduced, so that the total load does not change.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 134) in this Chapter for the definition of self-weight for the Frame element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Distributed Span Load

The Distributed Span Load is used to apply distributed forces on Cable elements. The load intensity may be specified as uniform or trapezoidal. However, the load is actually applied as a uniform load per unit of undeformed length of the cable.

The total load is calculated and divided by the undeformed length to determine the magnitude of load to apply. As the cable stretches, the magnitude is correspondingly reduced, so that the total load does not change.

The direction of loading may be specified in a fixed coordinate system (global or alternate coordinates) or in the element local coordinate system.

To model the effect of a non-uniform distributed load on a catenary cable object, specify multiple segments for the single cable object. The distributed load on the object will be applied as piecewise uniform loads over the segments.

For more information:

- See Topic “Distributed Span Load” (page 137) in Chapter “The Frame Element.”
- See Chapter “Objects and Elements” (page 7) for how a single cable object is meshed into elements (segments) at analysis time.
- See Chapter “Load Patterns” (page 321).

Temperature Load

Temperature Load creates axial thermal strain in the Cable element. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the element. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis.

The Load Temperature may be constant along the element length or interpolated from values given at the joints.

See Chapter “Load Patterns” (page 321) for more information.

Strain and Deformation Load

Axial Strain and Deformation Load change the length of the cable element. Deformation Load is the total change in length, whereas Strain Load is the change in length per unit of undeformed length. Positive values of these loads increase sag and tend to reduce tension in the cable, while negative values tighten up the cable and tend to increase tension.

See Chapter “Load Patterns” (page 321) for more information.

Target-Force Load

Target-Force Load is a special type of loading where you specify a desired cable tension, and deformation load is iteratively applied to achieve the target tension.

Since the tension may vary along the length of the cable, you must also specify the relative location where the desired tension is to occur. Target-Force loading is only used for nonlinear static and staged-construction analysis. If applied in any other type of Load Case, it has no effect.

Unlike all other types of loading, target-force loading is not incremental. Rather, you are specifying the total force that you want to be present in the cable element at the end of the Load Case or construction stage. The applied deformation that is calculated to achieve that force may be positive, negative, or zero, depending on the force present in the element at the beginning of the analysis. When a scale factor is applied to a Load Pattern that contains Target-Force loads, the total target force is scaled. The increment of applied deformation that is required may change by a different scale factor.

See Topic “Target-Force Load” (page 331) in Chapter “Load Patterns” and Topic “Target-Force Iteration” (page 444) in Chapter “Nonlinear Static Analysis” for more information.

Nonlinear Analysis

Nonlinear analysis is required to get meaningful results with the Cable element. Linear analyses can be performed, but they should always use the stiffness from the end of a nonlinear static Load Case in which all cables are loaded by their self-weight or other transverse load. For cases where no transverse load is present on a slack Cable element, the program will internally assume a very small self-weight load in order to obtain a unique shape. However, it is better if you apply a realistic load for this purpose.

Models with Cable elements will usually converge better if you allow a large number of Newton-Raphson iterations in the Load Case, say 25 or more. Convergence behavior is generally improved by using fewer segments in the cable object, and by applying larger load increments. Note that this is the opposite behavior than can be expected for cables modeled as frames, where using more segments and smaller load increments is usually advantageous.

Element Output

The catenary Cable element produces axial force (tension only) and displacement output along its length.

The Shell Element

The **Shell** element is a type of area object that is used to model membrane, plate, and shell behavior in planar and three-dimensional structures. The shell material may be homogeneous or layered through the thickness. Material nonlinearity can be considered when using the layered shell.

Basic Topics for All Users

- Overview
- Joint Connectivity
- Edge Constraints
- Degrees of Freedom
- Local Coordinate System
- Section Properties
- Mass
- Self-Weight Load
- Uniform Load
- Surface Pressure Load
- Internal Force and Stress Output

Advanced Topics

- Advanced Local Coordinate System
- Property Modifiers
- Joint Offsets and Thickness Overwrites
- Gravity Load
- Temperature Load

Overview

The Shell element is a three- or four-node formulation that combines membrane and plate-bending behavior. The four-joint element does not have to be planar.

Each Shell element has its own local coordinate system for defining Material properties and loads, and for interpreting output. Temperature-dependent, orthotropic material properties are allowed. Each element may be loaded by gravity and uniform loads in any direction; surface pressure on the top, bottom, and side faces; and loads due to strain and temperature change.

A four-point numerical integration formulation is used for the Shell stiffness. Stresses and internal forces and moments, in the element local coordinate system, are evaluated at the 2-by-2 Gauss integration points and extrapolated to the joints of the element. An approximate error in the element stresses or internal forces can be estimated from the difference in values calculated from different elements attached to a common joint. This will give an indication of the accuracy of a given finite-element approximation and can then be used as the basis for the selection of a new and more accurate finite element mesh.

Structures that can be modeled with this element include:

- Floor systems
- Wall systems
- Bridge decks
- Three-dimensional curved shells, such as tanks and domes
- Detailed models of beams, columns, pipes, and other structural members

Two distinct formulations are available: homogenous and layered.

Homogeneous

The homogeneous shell combines independent membrane and plate behavior. These behaviors become coupled if the element is warped (non-planar.) The membrane behavior uses an isoparametric formulation that includes translational in-plane stiffness components and a “drilling” rotational stiffness component in the direction normal to the plane of the element. See Taylor and Simo (1985) and Ibrahimbegovic and Wilson (1991). In-plane displacements are quadratic.

Plate-bending behavior includes two-way, out-of-plane, plate rotational stiffness components and a translational stiffness component in the direction normal to the plane of the element. You may choose a thin-plate (Kirchhoff) formulation that neglects transverse shearing deformation, or a thick-plate (Mindlin/Reissner) formulation which includes the effects of transverse shearing deformation. Out-of-plane displacements are cubic.

For each homogeneous Shell element in the structure, you can choose to model pure-membrane, pure-plate, or full-shell behavior. It is generally recommended that you use the full shell behavior unless the entire structure is planar and is adequately restrained.

Layered

The layered shell allows any number of layers to be defined in the thickness direction, each with an independent location, thickness, behavior, and material. Material behavior may be nonlinear.

Membrane deformation within each layer uses a strain-projection method (Hughes, 2000.) In-plane displacements are quadratic. Unlike for the homogeneous shell, the “drilling” degrees of freedom are not used, and they should not be loaded. These rotations normal to the plane of the element are only loosely tied to the rigid-body rotation of the element to prevent instability.

For bending, a Mindlin/Reissner formulation is used which always includes transverse shear deformations. Out-of-plane displacements are quadratic and are consistent with the in-plane displacements.

The layered Shell usually represents full-shell behavior, although you can control this on a layer-by-layer basis. Unless the layering is fully symmetrical in the thickness direction, membrane and plate behavior will be coupled.

Joint Connectivity

Each Shell element (and other types of area objects/elements) may have either of the following shapes, as shown in Figure 42 (page 181):

- Quadrilateral, defined by the four joints **j1**, **j2**, **j3**, and **j4**.
- Triangular, defined by the three joints **j1**, **j2**, and **j3**.

The quadrilateral formulation is the more accurate of the two. The triangular element is only recommended for locations where the stresses do not change rapidly. The use of large triangular elements is not recommended where in-plane (membrane) bending is significant. The use of the quadrilateral element for meshing various geometries and transitions is illustrated in Figure 43 (page 182), so that triangular elements can be avoided altogether.

Edge constraints are also available to create transitions between mis-matched meshes without using distorted elements. See Subtopic “Edge Constraints” (page 183) for more information.

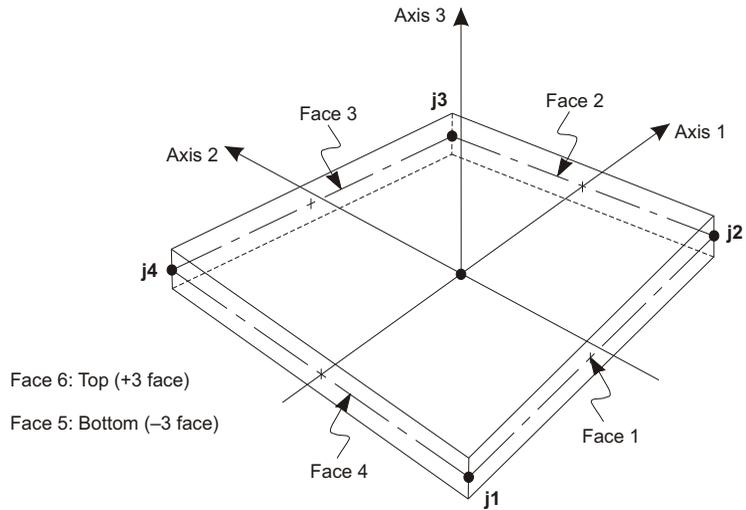
The joints **j1** to **j4** define the corners of the reference surface of the shell element. For the homogeneous shell this is the mid-surface of the element; for the layered shell you choose the location of this surface relative to the material layers.

You may optionally assign joint offsets to the element that shift the reference surface away from the joints. See Topic “Joint Offsets and Thickness Overwrites” (page 203) for more information.

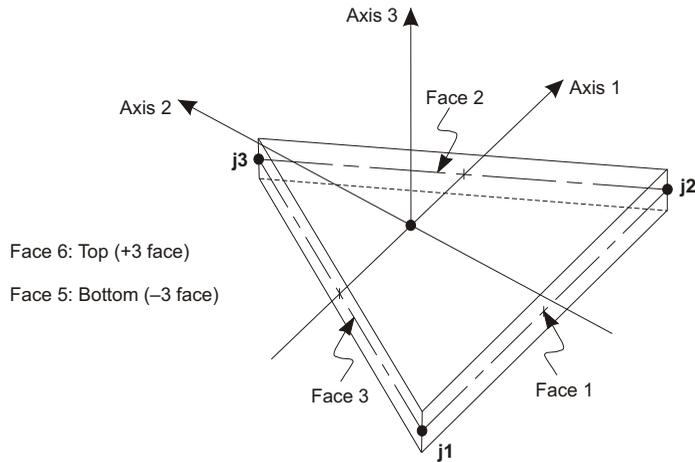
Shape Guidelines

The locations of the joints should be chosen to meet the following geometric conditions:

- The inside angle at each corner must be less than 180° . Best results for the quadrilateral will be obtained when these angles are near 90° , or at least in the range of 45° to 135° .
- The aspect ratio of an element should not be too large. For the triangle, this is the ratio of the longest side to the shortest side. For the quadrilateral, this is the ratio of the longer distance between the midpoints of opposite sides to the shorter such distance. Best results are obtained for aspect ratios near unity, or at least less than four. The aspect ratio should not exceed ten.



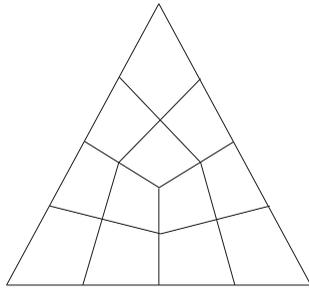
Four-node Quadrilateral Shell Element



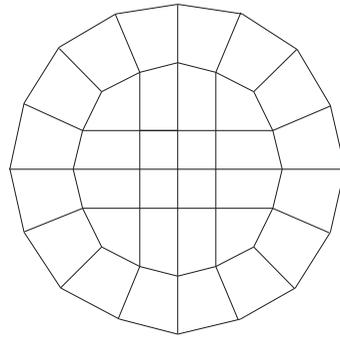
Three-node Triangular Shell Element

Figure 42
Area Element Joint Connectivity and Face Definitions

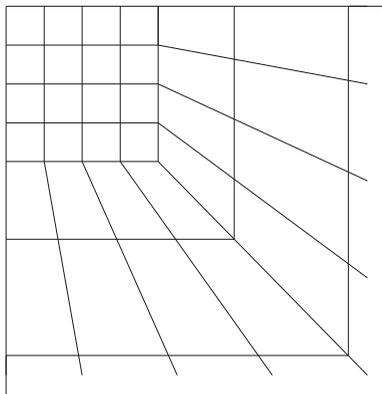
- For the quadrilateral, the four joints need not be coplanar. A small amount of twist in the element is accounted for by the program. The angle between the normals at the corners gives a measure of the degree of twist. The normal at a



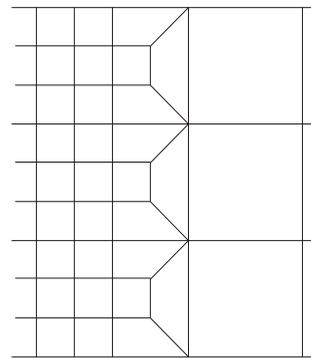
Triangular Region



Circular Region



Infinite Region



Mesh Transition

Figure 43
Mesh Examples Using the Quadrilateral Area Element

corner is perpendicular to the two sides that meet at the corner. Best results are obtained if the largest angle between any pair of corners is less than 30° . This angle should not exceed 45° .

These conditions can usually be met with adequate mesh refinement. The accuracy of the thick-plate and layered formulations is more sensitive to large aspect ratios and mesh distortion than is the thin-plate formulation.

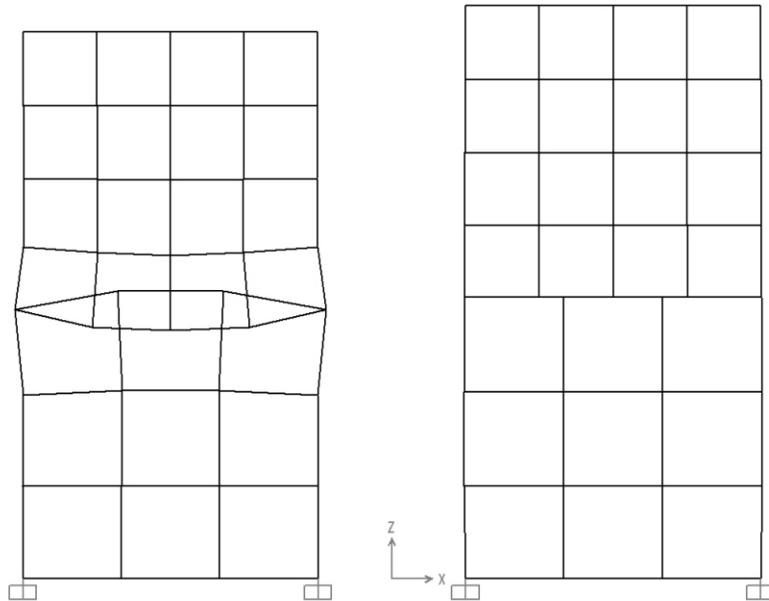


Figure 44

Connecting Meshes with the Edge Constraints: Left Model – No Edge Constraints; Right Model – Edge Constraints Assigned to All Elements

Edge Constraints

You can assign automatic **edge constraints** to any shell element (or any area objects.) When edge constraints are assigned to an element, the program automatically connects all joints that are on the edge of the element to the adjacent corner joints of the element. Joints are considered to be on the edge of the element if they fall within the auto-merge tolerance set by you in the Graphical User Interface.

Edge constraints can be used to connect together mis-matched shell meshes, but will also connect *any* element that has a joint on the edge of the shell to that shell. This include beams, columns, restrained joints, link supports, etc.

These joints are connected by flexible interpolation constraints. This means that the displacements at the intermediate joints on the edge are interpolated from the displacements of the corner joints of the shell. No overall stiffness is added to the model; the effect is entirely local to the edge of the element.

Figure 44 (page 183) shows an example of two mis-matched meshes, one connected with edge constraints, and one not. In the connected mesh on the right, edge constraints were assigned to all elements, although it was really only necessary to do so for the elements at the transition. Assigning edge constraints to elements that do not need them has little effect on performance and no effect on results.

The advantage of using edge constraints instead of the mesh transitions shown in Figure 43 (page 182) is that edge constraints do not require you to create distorted elements. This can increase the accuracy of the results.

It is important to understand that near any transition, whether using edge constraints or not, the accuracy of stress results is controlled by the largest element size. Furthermore, the effect of the coarser mesh propagates into the finer mesh for a distance that is on the order of the size of the larger elements, as governed by St. Venant's effect. For this reason, be sure to create your mesh transitions far enough away from the areas where you need detailed stress results.

Important Note: Edge constraints transfer load from intermediate joints to corner joints along the edge of a shell. Applying an edge constraint along an edge that is co-linear with a frame, cable, tendon or link object can result in load being transferred by the edge constraint instead of by that object. Using edge constraints in these locations should be avoided if detailed results in the frame/cable/tendon/link are of interest. In particular, frame design results could be affected, and may be unconservative. Frame/cable/tendon/link objects with only one joint connected to the edge are not affected, and in fact one of the advantages of using edge constraints is to connect such elements to a coarse shell mesh, provided that detailed local stresses in the shells are not needed. When frame/cable/tendon/link objects are co-linear with an edge constraint, the overall effect of the object on the model is captured, but local response may not be accurate.

Degrees of Freedom

The Shell element always activates all six degrees of freedom at each of its connected joints. When the element is used as a pure membrane, you must ensure that restraints or other supports are provided to the degrees of freedom for normal translation and bending rotations. When the element is used as a pure plate, you must ensure that restraints or other supports are provided to the degrees of freedom for in-plane translations and the rotation about the normal.

The use of the full shell behavior (membrane plus plate) is recommended for all three-dimensional structures.

Note that the “drilling” degree of freedom (rotation about the normal) is not used for the layered shell and should not be loaded.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom” for more information.

Local Coordinate System

Each Shell element (and other types of area objects/elements) has its own **element local coordinate system** used to define Material properties, loads and output. The axes of this local system are denoted 1, 2 and 3. The first two axes lie in the plane of the element with an orientation that you specify; the third axis is normal.

It is important that you clearly understand the definition of the element local 1-2-3 coordinate system and its relationship to the global X-Y-Z coordinate system. Both systems are right-handed coordinate systems. It is up to you to define local systems which simplify data input and interpretation of results.

In most structures the definition of the element local coordinate system is extremely simple. The methods provided, however, provide sufficient power and flexibility to describe the orientation of Shell elements in the most complicated situations.

Local axes are first computed for the element without considering joint offsets. These are called the nominal local axes. If the joint offsets shift the reference surface by a different amount at each joint, the local axes are then transformed by projecting them onto the new normal to determine the actual local coordinate system used for analysis.

The discussion below considers the calculation of the nominal local axes using the joints. The transformation for the joint offsets, if needed, is discussed later in Topic “Joint Offsets and Thickness Overwrites” (page 203).

The simplest method, using the **default orientation** and the **Shell element coordinate angle**, is described in this topic. Additional methods for defining the Shell element local coordinate system are described in the next topic.

For more information:

- See Chapter “Coordinate Systems” (page 11) for a description of the concepts and terminology used in this topic.
- See Topic “Advanced Local Coordinate System” (page 186) in this Chapter.

Normal Axis 3

Local axis 3 is always normal to the plane of the Shell element. This axis is directed toward you when the path **j1-j2-j3** appears counterclockwise. For quadrilateral elements, the element plane is defined by the vectors that connect the midpoints of the two pairs of opposite sides.

Default Orientation

The default orientation of the local 1 and 2 axes is determined by the relationship between the local 3 axis and the global Z axis:

- The local 3-2 plane is taken to be vertical, i.e., parallel to the Z axis
- The local 2 axis is taken to have an upward (+Z) sense unless the element is horizontal, in which case the local 2 axis is taken along the global +Y direction
- The local 1 axis is horizontal, i.e., it lies in the X-Y plane

The element is considered to be horizontal if the sine of the angle between the local 3 axis and the Z axis is less than 10^{-3} .

The local 2 axis makes the same angle with the vertical axis as the local 3 axis makes with the horizontal plane. This means that the local 2 axis points vertically upward for vertical elements.

Element Coordinate Angle

The Shell element coordinate angle, **ang**, is used to define element orientations that are different from the default orientation. It is the angle through which the local 1 and 2 axes are rotated about the positive local 3 axis from the default orientation. The rotation for a positive value of **ang** appears counterclockwise when the local +3 axis is pointing toward you.

For horizontal elements, **ang** is the angle between the local 2 axis and the horizontal +Y axis. Otherwise, **ang** is the angle between the local 2 axis and the vertical plane containing the local 3 axis. See Figure 45 (page 187) for examples.

Advanced Local Coordinate System

By default, the element local coordinate system is defined using the element coordinate angle measured with respect to the global +Z and +Y directions, as described

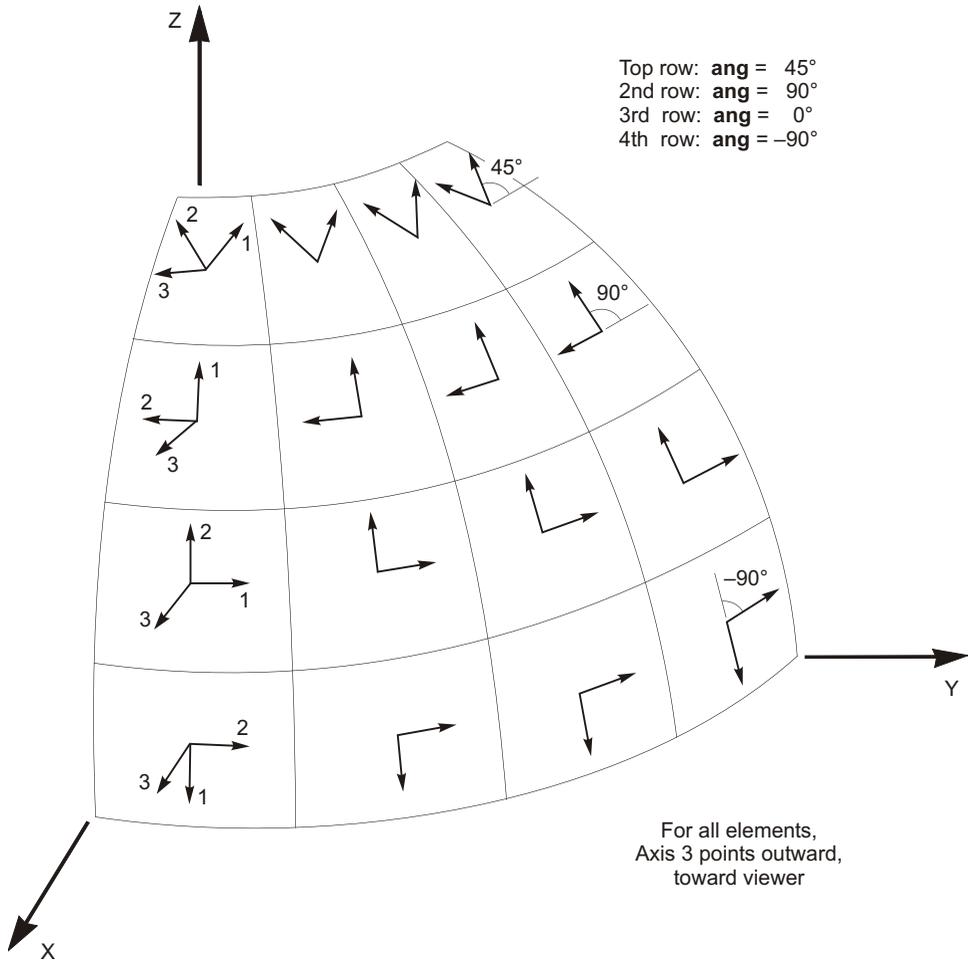


Figure 45
The Area Element Coordinate Angle with Respect to the Default Orientation

in the previous topic. In certain modeling situations it may be useful to have more control over the specification of the local coordinate system.

This topic describes how to define the orientation of the tangential local 1 and 2 axes, with respect to an arbitrary reference vector when the element coordinate angle, **ang**, is zero. If **ang** is different from zero, it is the angle through which the local 1 and 2 axes are rotated about the positive local 3 axis from the orientation determined by the reference vector. The local 3 axis is always normal to the plane of the element.

For more information:

- See Chapter “Coordinate Systems” (page 11) for a description of the concepts and terminology used in this topic.
- See Topic “Local Coordinate System” (page 185) in this Chapter.

Reference Vector

To define the tangential local axes, you specify a **reference vector** that is parallel to the desired 3-1 or 3-2 plane. The reference vector must have a positive projection upon the corresponding tangential local axis (1 or 2, respectively). This means that the positive direction of the reference vector must make an angle of less than 90 with the positive direction of the desired tangential axis.

To define the reference vector, you must first specify or use the default values for:

- A primary coordinate direction **pldirp** (the default is +Z)
- A secondary coordinate direction **pldirs** (the default is +Y). Directions **pldirs** and **pldirp** should not be parallel to each other unless you are sure that they are not parallel to local axis 3
- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system)
- The local plane, **local**, to be determined by the reference vector (the default is 32, indicating plane 3-2)

You may optionally specify:

- A pair of joints, **plveca** and **plvecb** (the default for each is zero, indicating the center of the element). If both are zero, this option is not used

For each element, the reference vector is determined as follows:

1. A vector is found from joint **plveca** to joint **plvecb**. If this vector is of finite length and is not parallel to local axis 3, it is used as the reference vector \mathbf{V}_p
2. Otherwise, the primary coordinate direction **pldirp** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis 3, it is used as the reference vector \mathbf{V}_p
3. Otherwise, the secondary coordinate direction **pldirs** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis 3, it is used as the reference vector \mathbf{V}_p

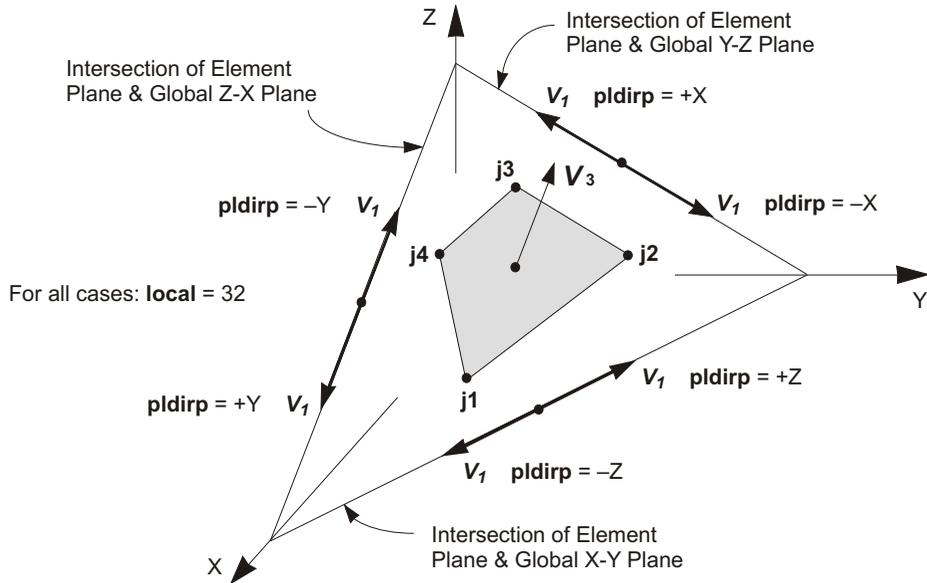


Figure 46
Area Element Local Coordinate System Using Coordinate Directions

4. Otherwise, the method fails and the analysis terminates. This will never happen if **pdirp** is not parallel to **pdirp**

A vector is considered to be parallel to local axis 3 if the sine of the angle between them is less than 10^{-3} .

The use of the coordinate direction method is illustrated in Figure 46 (page 189) for the case where **local** = 32.

Determining Tangential Axes 1 and 2

The program uses vector cross products to determine the tangential axes 1 and 2 once the reference vector has been specified. The three axes are represented by the three unit vectors \mathbf{V}_1 , \mathbf{V}_2 and \mathbf{V}_3 , respectively. The vectors satisfy the cross-product relationship:

$$\mathbf{V}_1 \times \mathbf{V}_2 = \mathbf{V}_3$$

The tangential axes 1 and 2 are defined as follows:

- If the reference vector is parallel to the 3-1 plane, then:

$$\begin{matrix} \mathbf{V}_2 & \mathbf{V}_3 & \mathbf{V}_p & \text{and} \\ \mathbf{V}_1 & \mathbf{V}_2 & \mathbf{V}_3 \end{matrix}$$

- If the reference vector is parallel to the 3-2 plane, then:

$$\begin{matrix} \mathbf{V}_1 & \mathbf{V}_p & \mathbf{V}_3 & \text{and} \\ \mathbf{V}_2 & \mathbf{V}_3 & \mathbf{V}_1 \end{matrix}$$

In the common case where the reference vector is parallel to the plane of the element, the tangential axis in the selected local plane will be equal to \mathbf{V}_p .

Section Properties

A **Shell Section** is a set of material and geometric properties that describe the cross-section of one or more Shell objects (elements.) A Shell Section property is a type of Area Section property. Sections are defined independently of the objects, and are assigned to the area objects.

Area Section Type

When defining an area section, you have a choice of three basic element types:

- Shell – the subject of this Chapter, with translational and rotational degrees of freedom, capable of supporting forces and moments
- Plane (stress or strain) – a two-dimensional solid, with translational degrees of freedom, capable of supporting forces but not moments. This element is covered in Chapter “The Plane Element” (page 215).
- Asolid – axisymmetric solid, with translational degrees of freedom, capable of supporting forces but not moments. This element is covered in Chapter “The Asolid Element” (page 225).

Shell Section Type

For Shell sections, you may choose one of the following types of behavior:

- Membrane
 - Pure membrane behavior
 - Supports only the in-plane forces and the normal (drilling) moment
 - Linear, homogeneous material.

- Plate
 - Pure plate behavior
 - Supports only the bending moments and the transverse force
 - Thick- or thin-plate formulation
 - Linear, homogeneous material.
- Shell
 - Full shell behavior, a combination of membrane and plate behavior
 - Supports all forces and moments
 - Thick- or thin-plate formulation
 - Linear, homogeneous material.
- Layered
 - Multiple layers, each with a different material, thickness, behavior, and location
 - Provides full-shell behavior unless all layers have only membrane or only plate behavior
 - With full-shell behavior, supports all forces and moments except the “drilling” moment
 - Thick-plate formulation; may be nonlinear.

It is generally recommended that you use the full-shell behavior unless the entire structure is planar and is adequately restrained.

Homogeneous Section Properties

Homogeneous material properties are used for the non-layered Membrane, Plate, and Shell section types. The following data needs to be specified.

Section Thickness

Each homogeneous Section has a constant membrane thickness and a constant bending thickness. The membrane thickness, **th**, is used for calculating:

- The membrane stiffness for full-shell and pure-membrane Sections
- The element volume for the element self-weight and mass calculations

The bending thickness, **thb**, is use for calculating:

- The plate-bending and transverse-shearing stiffnesses for full-shell and pure-plate Sections

Normally these two thicknesses are the same and you only need to specify **th**. However, for some applications, you may wish to artificially change the membrane or plate stiffness. For this purpose, you may specify a value of **thb** that is different from **th**. For more detailed control, such as representing corrugated or orthotropic construction, the use of property modifiers is better. See Topic “Property Modifiers” (page 201.)

Thickness Formulation

Two thickness formulations are available, which determine whether or not transverse shearing deformations are included in the plate-bending behavior of a plate or shell element:

- The thick-plate (Mindlin/Reissner) formulation, which includes the effects of transverse shear deformation
- The thin-plate (Kirchhoff) formulation, which neglects transverse shearing deformation

Shearing deformations tend to be important when the thickness is greater than about one-tenth to one-fifth of the span. They can also be quite significant in the vicinity of bending-stress concentrations, such as near sudden changes in thickness or support conditions, and near holes or re-entrant corners.

Even for thin-plate bending problems where shearing deformations are truly negligible, the thick-plate formulation tends to be more accurate, although somewhat stiffer, than the thin-plate formulation. However, the accuracy of the thick-plate formulation is more sensitive to large aspect ratios and mesh distortion than is the thin-plate formulation.

It is generally recommended that you use the thick-plate formulation unless you are using a distorted mesh and you know that shearing deformations will be small, or unless you are trying to match a theoretical thin-plate solution.

The thickness formulation has no effect upon membrane behavior, only upon plate-bending behavior.

Section Material

The material properties for each Section are specified by reference to a previously-defined Material. The material may be isotropic, uniaxial, or orthotropic. If

an anisotropic material is chosen, orthotropic properties will be used. The material properties used by the Shell Section are:

- The moduli of elasticity, **e1**, **e2**, and **e3**
- The shear modulus, **g12**, **g13**, and **g23**
- The Poisson's ratios, **u12**, **u13**, and **u23**
- The coefficients of thermal expansion, **a1** and **a2**
- The mass density, **m**, for computing element mass
- The weight density, **w**, for computing Self-Weight and Gravity Loads

The properties **e3**, **u13**, and **u23** are condensed out of the material matrix by assuming a state of plane stress in the element. The resulting, modified values of **e1**, **e2**, **g12**, and **u12** are used to compute the membrane and plate-bending stiffnesses.

The shear moduli, **g13** and **g23**, are used to compute the transverse shearing stiffness if the thick-plate formulation is used. The coefficients of thermal expansion, **a1** and **a2**, are used for membrane expansion and thermal bending strain.

All material properties (except the densities) are obtained at the material temperature of each individual element.

See Chapter “Material Properties” (page 69) for more information.

Section Material Angle

The material local coordinate system and the element (Shell Section) local coordinate system need not be the same. The local 3 directions always coincide for the two systems, but the material 1 axis and the element 1 axis may differ by the angle **a** as shown in Figure 47 (page 194). This angle has no effect for isotropic material properties since they are independent of orientation.

See Topic “Local Coordinate System” (page 70) in Chapter “Material Properties” for more information.

Layered Section Property

For the layered Section property, you define how the section is built-up in the thickness direction. Any number of layers is allowed, even a single layer. Layers are located with respect to a reference surface. This reference surface may be the middle surface, the neutral surface, the top, the bottom, or any other location you choose. By default, the reference surface contains the element nodes, although this can be changed using joint offsets.

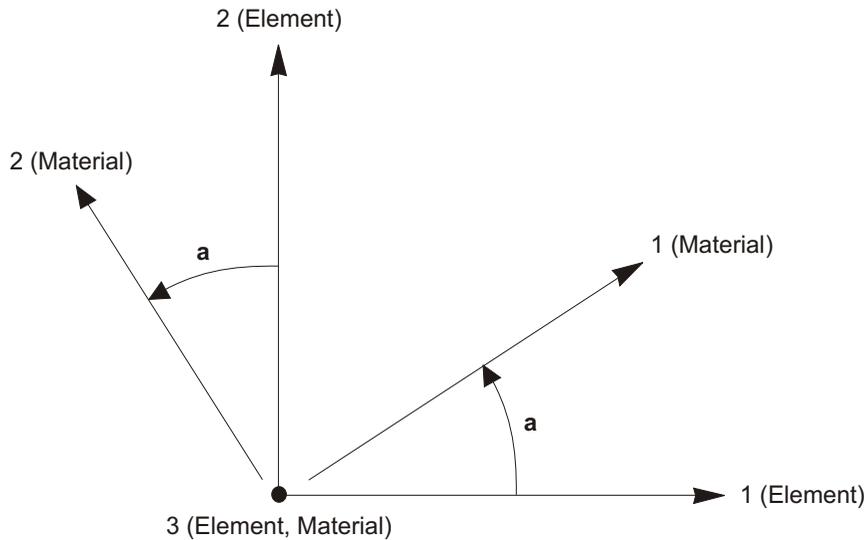


Figure 47
Shell Section Material Angle

The thick-plate (Mindlin/Reissner) formulation, which includes the effects of transverse shear deformation, is always used for bending behavior the layered shell.

The following eight parameters are specified to define each layer, as illustrated in Figure 48 (page 195.)

(1) Layer Name

The layer name is arbitrary, but must be unique within a single Section. However, the same layer name can be used in different Sections. This can be useful because results for a given layer name can be plotted simultaneously for elements having different Sections.

(2) Layer Distance

Each layer is located by specifying the distance from the reference surface to the center of the layer, measured in the positive local-3 direction of the element. This value is called **d** in the examples below.

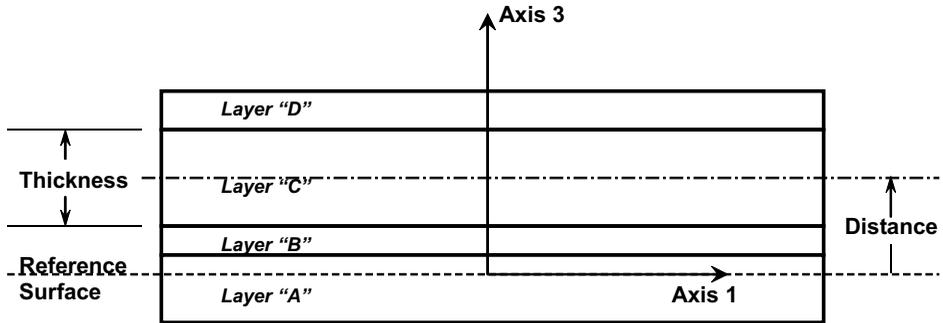


Figure 48

Four-Layer Shell, Showing the Reference Surface, the Names of the Layers, and the Distance and Thickness for Layer "C"

(3) Layer Thickness

Each layer has a single thickness, measured in the local-3 direction of the element. For modeling rebar or material fibers, you can specify a very thin "smeared" layer that has an equivalent cross-sectional area. This value is called **th** in the examples below.

(4) Layer Type

You can choose between:

- Membrane: Strains in the layer $(\epsilon_{11}, \epsilon_{22}, \epsilon_{12})$ are computed only from in-plane membrane displacements, and stresses in the layer $(\sigma_{11}, \sigma_{22}, \sigma_{12})$ contribute only to in-plane membrane forces (F_{11}, F_{22}, F_{12}) .
- Plate: Strains in the layer $(\epsilon_{11}, \epsilon_{22}, \epsilon_{12}, \epsilon_{13}, \epsilon_{23})$ are computed only from plate-bending rotations and transverse displacements, and stresses in the layer $(\sigma_{11}, \sigma_{22}, \sigma_{12}, \sigma_{13}, \sigma_{23})$ contribute only to plate-bending moments and transverse shearing forces $(M_{11}, M_{22}, M_{12}, V_{13}, V_{23})$.
- Shell, which combines membrane and plate behavior: Strains in the layer $(\epsilon_{11}, \epsilon_{22}, \epsilon_{12}, \epsilon_{13}, \epsilon_{23})$ are computed from all displacements and plate-bending rotations, and stresses in the layer $(\sigma_{11}, \sigma_{22}, \sigma_{12}, \sigma_{13}, \sigma_{23})$ contribute to

all forces and plate-bending moments
($F_{11}, F_{22}, F_{12}, M_{11}, M_{22}, M_{12}, V_{13}, V_{23}$).

In most applications, layers should use shell behavior. See shear-wall modeling below for an example of where you might want to separate membrane and plate behavior.

Important Note: Mass and weight are computed only for membrane and shell layers, not for plate layers. This prevents double-counting when independent membrane and plate layers are used for the same material.

(5) Layer Number of Thickness Integration Points

Material behavior is integrated (sampled) at a finite number of points in the thickness direction of each layer. You may choose one to five points for each layer. The location of these points follows standard Gauss integration procedures. This value is called **n** in the examples below.

For a single layer of linear material, one point in the thickness direction is adequate to represent membrane behavior, and two points will capture both membrane and plate behavior. If you have multiple layers, you may be able to use a single point for thinner layers.

Nonlinear behavior may require more integration points or more layers in order to capture yielding near the top and bottom surfaces. Using an excessive number of integration points can increase analysis time. You may need to experiment to find a balance between accuracy and computational efficiency.

(6) Layer Material

The material properties for each layer are specified by reference to a previously-defined Material. The material may be isotropic, uniaxial, or orthotropic. If an anisotropic material is chosen, orthotropic properties will be used. The behavior of the material depends on the material component behavior chosen for the layer, as described below.

(7) Layer Material Angle

For orthotropic and uniaxial materials, the material axes may be rotated with respect to the element axes. Each layer may have a different material angle. For example, you can model rebar in two orthogonal directions as two layers of uniaxial material with material angles 90° apart. This value is called **ang** in the examples below. For further information, see topic “Section Material Angle” above (page 193.)

(8) Layer Material Behavior

Choose between “Directional” and “Coupled”. Directional behavior can be applied to all materials and is described in the remainder of this topic. Coupled behavior is available for concrete materials only, and uses the modified Darwin-Pecknold behavior as described in Chapter “Material Properties”, Topic “Modified Darwin-Pecknold Concrete Model” (page 100).

(9) Layer Material Components

This option applies only to “Directional” material behavior. For each of the three membrane stress components (σ_{11} , σ_{22} , σ_{12}), you can choose whether the behavior is linear, nonlinear, or inactive. For a uniaxial material, only the two components (σ_{11} , σ_{12}) are significant, since $\sigma_{22} = 0$ always. Material components are defined in the material local coordinate system, which depends on the material angle and may not be the same for every layer.

If all three components are linear (two for the uniaxial material), then the linear material matrix is used for the layer, according to Equations (1) to (4) in Chapter “Material Properties” (page 69). Note that for anisotropic materials, the shear coupling terms in Equation (4) are neglected so that the behavior is the same as given by Equation (3).

If one or more of the three components is nonlinear or inactive, then all linear components use an uncoupled isotropic linear stress-strain law, all nonlinear components use the nonlinear stress-strain relationship, and all inactive components assume zero stress. The components become uncoupled, and behave as if Poisson’s ratio is zero. The behavior is summarized in the following table:

Component	Linear	Nonlinear	Inactive
σ_{11}	$\sigma_{11} \mathbf{e1} \sigma_{11}$	Eqns. (5)	$\sigma_{11} \quad 0$
σ_{22}	$\sigma_{22} \mathbf{e1} \sigma_{22}$	Eqns. (5)	$\sigma_{22} \quad 0$
σ_{12}	$\sigma_{12} \mathbf{e1} \sigma_{12}$	Eqns. (6)	$\sigma_{12} \quad 0$

Note that the linear equation for σ_{12} is for an isotropic material with zero Poisson’s ratio. See Chapter “Material Properties” (page 69) for Equations (5) and (6).

For a uniaxial material, $\sigma_{22} = 0$ and σ_{12} is half the value given in the table above.

Transverse shear behavior is always linear, and is controlled by the corresponding moment components. For a layer of type Membrane, the transverse shear stresses (τ_{13} , τ_{23}) are both zero. For a layer of type Plate or Shell:

- $\tau_{13} = 0$ if M_{11} is inactive, else $\tau_{13} = \mathbf{g13} \cdot M_{13}$
- $\tau_{23} = 0$ if M_{22} is inactive, else $\tau_{23} = \mathbf{g23} \cdot M_{23}$

Interaction Between Layers

Layers are defined independently, and it is permissible for layers to overlap, or for gaps to exist between the layers. It is up to you to decide what is appropriate.

For example, when modeling a concrete slab, you can choose a single layer to represent the full thickness of concrete, and four layers to represent rebar (two near the top at a 90° angle to each other, and two similar layers at the bottom.) These rebar layers would be very thin, using an equivalent thickness to represent the cross-sectional area of the steel. Because the layers are so thin, there is no need to worry about the fact that the rebar layers overlap the concrete. The amount of excess concrete that is contained in the overlapped region is very small.

Layers are kinematically connected by the Mindlin/Reissner assumption that normals to the reference surface remain straight after deformation. This is the shell equivalent to the beam assumption that plane sections remain plane.

Integration in the Plane

Force-deflection behavior is computed by integrating the stress-strain behavior through the thickness and over the 1-2 plane of the element. You can specify the number of integration points in the thickness direction of each layer as described above.

For each of these thickness locations, integration in the plane is performed at the standard 2 x 2 Gauss points (coordinates 0.577 on a square of size 1.0). Nonlinear behavior is sampled only at these points. This is equivalent to having two fibers, located approximately at the 1/4 and 3/4 points, in each of the local 1 and 2 directions. Plotted or tabulated stresses at locations other than the four Gauss points are interpolated or extrapolated, and do not necessarily represent the sampled nonlinear stresses. For this reason, stresses at the joints may sometimes appear to exceed failure stresses.

Example: Nonlinear Shear-Wall, “Realistic” Modeling

An important application for the layered shell element is nonlinear shear-wall modeling, and it will serve as an example for other applications. Let's consider an 18 inch (457 mm) thick vertical wall, with two vertical and two horizontal layers of rebar having 3 inch (76 mm) cover from both faces. The two horizontal layers together provide a 1% rebar area ratio, and the two vertical layers together provide an area ratio of 2%.

When modeling linear behavior, it is not usually necessary to include the rebar, but it is essential for nonlinear behavior. In the simplest case, the entire wall section will be considered as nonlinear for both membrane and bending behavior, leading to the most “realistic”, if not the most practical model. This requires a layered section with five layers:

“Realistic” Shear-Wall Model

Layer	Type	Material	th	d	ang	n	11	22	12
1	Shell	Conc	18.00	0.	0	5	N	N	N
2	Shell	Rebar	0.09	+6.	0	1	N		N
3	Shell	Rebar	0.09	6.	0	1	N		N
4	Shell	Rebar	0.18	+6.	90	1	N		N
5	Shell	Rebar	0.18	6.	90	1	N		N

For the stress components, “N” indicates nonlinear, “L” indicates linear, and “ ” indicates inactive. The concrete material behavior may be “Directional” or “Coupled”.

Note that for the rebar, σ_{11} is always nonlinear. Vertical rebar is defined by setting the material angle to 90°, which aligns it with the shell local-2 axis. Hence the vertical rebar stress σ_{11} corresponds to shell σ_{22} .

Also note that for the rebar, σ_{12} is set to be nonlinear. This allows the rebar to carry shear when the concrete cracks. This can be taken to represent dowel action, although no information on actual dowel behavior is present in the model, so it is only an approximation. You must use your engineering judgement to determine if this approach is suitable to your needs. The most conservative approach is to set the rebar stress component σ_{12} to be inactive.

Example: Nonlinear Shear-Wall, “Practical” Modeling

The five-layer model above seems realistic, but presents many failure mechanisms that may cloud the engineering information required for performance-based design. Whenever possible, the simplest model should be used to meet the engineering goals. Doing this will make the analysis run faster and make the interpretation of results easier.

With this in mind, a more practical model is presented below, with only the vertical membrane stresses taken to be nonlinear. Such a model may be suitable for taller shear walls where column-like behavior governs:

“Practical” Shear-Wall Model

Layer	Type	Material	th	d	ang	n	11	22	12
1	Membr	Conc	18.00	0.	0	1	L	N	L
2	Membr	Rebar	0.18	+6.	90	1	N		
3	Membr	Rebar	0.18	6.	90	1	N		
4	Plate	Conc	16.00	0.	0	2	L	L	L

In this model, only membrane behavior is nonlinear, and only for the vertical stress component σ_{22} . This corresponds to rebar stress component σ_{11} when the material angle is 90° . Concrete material behavior is “Directional”.

It is generally not necessary to include rebar for linear behavior, so the horizontal rebar is omitted, and the rebar shear stress component τ_{12} is set to be inactive.

Out-of-plane behavior is assumed to be linear, so a single concrete plate layer is used. The thickness has been reduced to account for cracking without explicit nonlinear modeling. Plate bending stiffness is proportional to the cube of the thickness.

Example: In-fill Panel

There are many ways to model an infill panel. Two approaches will be presented here, both intended to represent membrane shear resistance only. The simplest is a single layer of concrete material carrying only membrane shear stress, as shown in the following model:

Infill Wall Simple Shear Model

Layer	Type	Material	th	d	ang	n	11	22	12
1	Membr	Conc	18.00	0.	0	1			N

In the second model, the concrete is assumed to act as compression struts along the two diagonals. For a square panel, these two struts would act at material angles of 45° , as shown in the following model:

Infill Wall Compression Strut Model

Layer	Type	Material	th	d	ang	n	11	22	12
1	Membr	Conc	18.00	0.	45	1	N		
2	Membr	Conc	18.00	0.	45	1	N		

Other possibilities exist. For both models, there is no vertical or horizontal membrane stiffness, and no plate-bending stiffness. Therefore, these models should only be used when the element is completely surrounded by frame or other supporting elements, and the elements should not be meshed.

Summary

As these examples show, you have considerable flexibility to create layered shell sections to represent a variety of linear and nonlinear behavior. The simplest model that accomplishes the engineering goals should be used. Even when more complicated models may be warranted, it is recommended to start with simple, mostly linear models, and increase the level of complexity and nonlinearity as you gain experience with your model and its behavior.

To assure a stable model, be sure to include layers that, when combined, provide both membrane and plate contributions to each of the three stress components.

Property Modifiers

You may specify scale factors to modify the computed section properties. These may be used, for example, to account for cracking of concrete, corrugated or orthotropic fabrication, or for other factors not easily described in the geometry and

material property values. Individual modifiers are available for the following ten terms:

- Membrane stiffness corresponding to force F11
- Membrane stiffness corresponding to force F22
- Membrane stiffness corresponding to force F12
- Plate bending stiffness corresponding to moment M11
- Plate bending stiffness corresponding to moment M22
- Plate bending stiffness corresponding to moment M12
- Plate shear stiffness corresponding to force V12
- Plate shear stiffness corresponding to force V13
- Mass
- Weight

The stiffness modifiers affect only homogenous elements, not layered elements. The mass and weight modifiers affect all elements.

See Topic “Internal Force and Stress Output” (page 210) for the definition of the force and moment components above.

You may specify multiplicative factors in two places:

- As part of the definition of the section property
- As an assignment to individual elements.

If modifiers are assigned to an element and also to the section property used by that element, then both sets of factors multiply the section properties.

Named Property Sets

In addition to directly assigning property modifiers to shell elements, you can apply them to a shell element in a staged-construction Load Case using a Named Property Set of Shell Property Modifiers. A Named Property Set includes the same ten factors above that can be assigned to an element.

When a Named Property Set is applied to an element in a particular stage of a Load Case, it replaces only the values that are assigned to the element or that had been applied in a previous stage; values computed by the Direct Analysis Method of design are also replaced. However, property modifiers specified with the section property remain in force and are not affected by the application of a Named Property Set.

The net effect is to use the factors specified in the Named Property Set multiplied by the factors specified in the section property.

When property modifiers are changed in a staged construction Load Case, they do not change the response of the structure up to that stage, but only affect subsequent response. In other words, the effect is incremental. For example, consider a cantilever with only default (unity) property modifiers, and a staged construction case as follows:

- Stage 1: Self-weight load is applied, resulting in a tip deflection of 1.0 and a support moment of 1000.
- Stage 2: Named Property Set “A” is applied that multiplies all stiffnesses by 2.0, and the mass and weight by 1.0. The tip deflection and support moment do not change.
- Stage 3: Self-weight load is applied again (incrementally). The resulting tip deflection is 1.5 and the support moment is 2000. Compared to Stage 1, the same incremental load is applied, but the structure is twice as stiff.
- Stage 4: Named Property Set “B” is applied that multiplies all stiffnesses, as well as the mass and weight, by 2.0. The tip deflection and support moment do not change.
- Stage 5: Self-weight load is applied again (incrementally). The resulting tip deflection is 2.5 and the support moment is 4000. Compared to Stage 1, twice the incremental load is applied, and the structure is twice as stiff

Joint Offsets and Thickness Overwrites

You may optionally assign joint offset and thickness overwrites to any element. These are often used together to align the top or bottom of the shell element with a given surface. See Figure 49 (page 204.)

Joint Offsets

Joint offsets are measured from the joint to the reference surface of the element in the direction normal to the plane of the joints. If the joints define a warped surface, the plane is determined by the two lines connecting opposite mid-sides (i.e., the middle of **j1-j2** to the middle of **j3-j4**, and the middle of **j1-j3** to the middle of **j2-j4**.) A positive offset is in the same direction as the positive local-3 axis of the element.

Joint offsets locate the reference plane of the element. For homogeneous shells, this is the mid-surface of the element. For layered shells, the reference surface is the

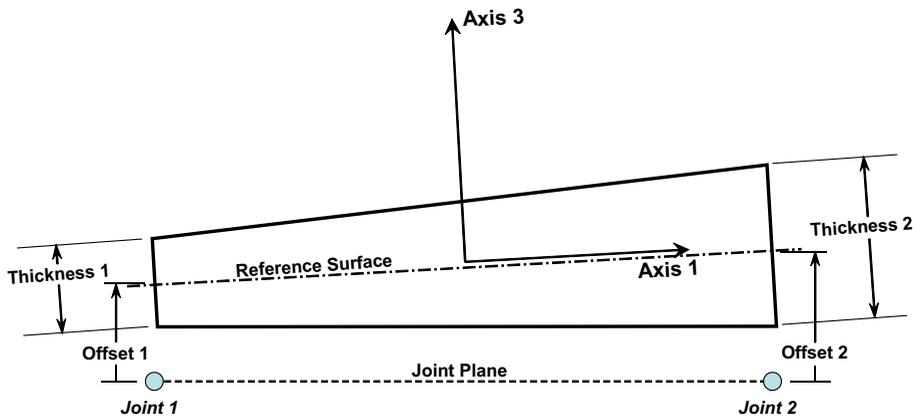


Figure 49
*Joint Offsets and Thickness Overwrites for a Homogeneous Shell
 Edge View shown Along One Side*

surface you used to locate the layers in the section. By changing the reference surface in a layered section, you can accomplish the same effect as using joint offsets that are equal at the joints. See Topic “Layered Section Property” (page 193) for more information.

When you assign joint offsets to a shell element, you can explicitly specify the offsets at the element joints, or you can reference a Joint Pattern. Using a Joint Pattern makes it easy to specify consistently varying offsets over many elements. See Topic “Joint Patterns” (page 332) in Chapter “Load Patterns” for more information.

Note that when the neutral surface of the element, after applying joint offsets, is no longer in the plane of the joints, membrane and plate-bending behavior become coupled. If you apply a diaphragm constraint to the joints, this will also constrain bending. Likewise, a plate constraint will constrain membrane action.

Effect of Joint Offsets on the Local Axes

The insertion points can interact significantly with the element local coordinate system. As described previously, the nominal local axes are computed for zero joint offsets, such the reference surface connects to the joints.

If there are no joint offsets, or the joint offsets are equal at the joints, the offset reference plane is parallel to the original reference plane of the joints. In this case no

further transformation is necessary, and the element local axes are the same as the nominal local axes.

If the joint offsets are not equal, a new local 3 axis ($\tilde{\mathbf{V}}_3$) is computed as the normal to the plane determined from the two lines connecting the midsides of the reference surface *after* applying the offsets to the joints.

If the normal has changed direction, then the element local coordinate system is computed as follows:

$$\begin{array}{l} \tilde{\mathbf{V}}_1 \quad \mathbf{V}_2 \quad \tilde{\mathbf{V}}_3 \quad \text{and} \\ \tilde{\mathbf{V}}_2 \quad \tilde{\mathbf{V}}_3 \quad \tilde{\mathbf{V}}_1 \end{array}$$

where \mathbf{V}_1 , \mathbf{V}_2 , and \mathbf{V}_3 and the nominal local axes computed previously based on the joints; and $\tilde{\mathbf{V}}_1$, $\tilde{\mathbf{V}}_2$, and $\tilde{\mathbf{V}}_3$ are the transformed local axes used for analysis. If \mathbf{V}_2 is vertical, then $\tilde{\mathbf{V}}_1$ will always be horizontal. Note that the two systems are identical if $\tilde{\mathbf{V}}_3 = \mathbf{V}_3$, the usual case.

The nominal axes are used only for determining the direction of joint offsets. The transformed axes are used for all analysis purposes, including loading and results output.

Thickness Overwrites

Normally the thickness of the shell element is defined by the Section Property assigned to the element. You have the option to overwrite this thickness, including the ability to specify a thickness that varies over the element.

Currently this option only affects homogeneous shells. The thickness of layered shells is not changed. When thickness overwrites are assigned to a homogeneous shell, both the membrane thickness, **th**, and the bending thickness, **thb**, take the overwritten value.

When you assign thickness overwrites to a shell element, you can explicitly specify the thicknesses at the element joints, or you can reference a Joint Pattern. Using a Joint Pattern makes it easy to specify consistently varying thickness over many elements. See Topic “Joint Patterns” (page 332) in Chapter “Load Patterns” for more information.

As an example, suppose you have a variable thickness slab, and you want the top surface to lie in a single flat plane. Define a Joint Pattern that defines the thickness over the slab. Draw the elements so that the joints lie in the top plane. Assign thick-

ness overwrites to all the elements using the Joint Pattern with a scale factor of one, and assign the joint offsets using the same Joint Pattern, but with a scale factor of one-half (positive or negative, as needed).

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Shell element is lumped at the element joints. *No inertial effects are considered within the element itself.*

The total mass of the element is equal to the integral over the plane of the element of the mass density, **m**, multiplied by the membrane thickness, **th**, for homogeneous sections, and the sum of the masses of the individual layers for layered sections. Note that for layered shells, mass is computed only for membrane and shell layers, not for plate layers. The total mass may be scaled by the mass property modifier.

The total mass is apportioned to the joints in a manner that is proportional to the diagonal terms of the consistent mass matrix. See Cook, Malkus, and Plesha (1989) for more information. The total mass is applied to each of the three translational degrees of freedom: UX, UY, and UZ. No mass moments of inertia are computed for the rotational degrees of freedom.

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties”.
- See Topic “Property Modifiers” (page 201) in this chapter.
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Static and Dynamic Analysis” (page 341).

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For a Shell element, the self-weight is a force that is uniformly distributed over the plane of the element. The magnitude of the self-weight is equal to the weight density, **w**, multiplied by the membrane thickness, **th**, for homogeneous sections, and the sum of the weights of the individual layers for layered sections. Note that for layered shells, weight is computed only for membrane and shell layers, not for plate layers. The total weight may be scaled by the weight property modifier.

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of w .
- See Topic “Property Modifiers” (page 201) in this chapter.
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Shell element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 191) in this Chapter for the definition of self-weight for the Shell element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Uniform Load

Uniform Load is used to apply uniformly distributed forces to the midsurfaces of the Shell elements. The direction of the loading may be specified in a fixed coordinate system (global or Alternate Coordinates) or in the element local coordinate system.

Load intensities are given as forces per unit area. Load intensities specified in different coordinate systems are converted to the element local coordinate system and added together. The total force acting on the element in each local direction is given by the total load intensity in that direction multiplied by the area of the mid-surface. This force is apportioned to the joints of the element.

Forces given in fixed coordinates can optionally be specified to act on the projected area of the mid-surface, i.e., the area that can be seen along the direction of loading.

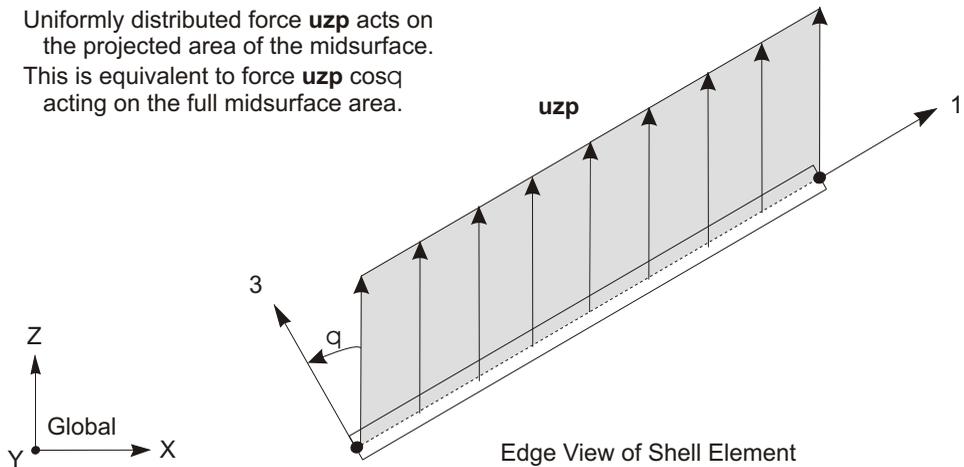


Figure 50

Example of Uniform Load Acting on the Projected Area of the Mid-surface

The specified load intensity is automatically multiplied by the cosine of the angle between the direction of loading and the normal to the element (the local 3 direction). This can be used, for example, to apply distributed snow or wind loads. See Figure 50 (page 208).

See Chapter “Load Patterns” (page 321) for more information.

Surface Pressure Load

The Surface Pressure Load is used to apply external pressure loads upon any of the six faces of the Shell element. The definition of these faces is shown in Figure 42 (page 181). Surface pressure always acts normal to the face. Positive pressures are directed toward the interior of the element.

The pressure may be constant over a face or interpolated from values given at the joints. The values given at the joints are obtained from Joint Patterns, and need not be the same for the different faces. Joint Patterns can be used to easily apply hydrostatic pressures.

The bottom and top faces are denoted Faces 5 and 6, respectively. The top face is the one visible when the +3 axis is directed toward you and the path **j1-j2-j3** ap-

pears counterclockwise. The pressure acting on the bottom or top face is integrated over the plane of the element and apportioned to the corner joints..

The sides of the element are denoted Faces 1 to 4 (1 to 3 for the triangle), counting counterclockwise from side **j1-j2** when viewed from the top. The pressure acting on a side is multiplied by the thickness, **th**, integrated along the length of the side, and apportioned to the two joints on that side. the bending thickness, **thb**, is not used.

For layered shells, the thickness used for edge loads is measured from the bottom of the bottom-most membrane or shell layer to the top of the top-most membrane or shell layer. Gaps between layers and overlapping layers do not change the thickness used. Plate layers are not considered when computing the loaded thickness.

For more information:

- See Topic “Thickness” (page 193) in this Chapter for the definition of **th**.
- See Chapter “Load Patterns” (page 321).

Temperature Load

Temperature Load creates thermal strain in the Shell element. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the element. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis.

Two independent Load Temperature fields may be specified:

- Temperature, **t**, which is constant through the thickness and produces membrane strains
- Temperature gradient, **t3**, which is linear in the thickness direction and produces bending strains

The temperature gradient is specified as the *change* in temperature per unit length. The temperature gradient is positive if the temperature increases (linearly) in the positive direction of the element local 3 axis. The gradient temperature is zero at the mid-surface, hence no membrane strain is induced.

Each of the two Load Temperature fields may be constant over the plane of the element or interpolated from values given at the joints.

See Chapter “Load Patterns” (page 321) for more information.

Strain Load

Eight types of strain load are available, one corresponding to each of the internal forces and moments in a shell element. These are:

- Membrane strain loads, ϵ_x , ϵ_y , and ϵ_z , representing change in the size and shape of the element that is uniform through the thickness. Positive membrane strain causes negative corresponding membrane forces in a restrained element.
- Bending strain loads, ϵ_{xy} , ϵ_{yz} , and ϵ_{zx} , representing change in the size and shape of the element that varies linearly through the thickness. Positive bending strain causes negative corresponding bending moments in a restrained element.
- Shear strain loads, ϵ_{13} and ϵ_{23} , representing the change in angle between the midsurface normal and the midsurface. Positive shear strain causes negative corresponding shear forces in a restrained element. Shear strain load has no effect on a thin shell or thin plate element, for which shear strain is assumed to be zero.

Any of the strain load fields may be constant over the plane of the element or interpolated from values given at the joints.

In an unrestrained element, strain loads cause deformation but induce no internal forces. This unrestrained deformation has the same sign as would deformation caused by the corresponding (conjugate) forces and moments acting on the element. On the other hand, strain loading in a restrained element causes corresponding internal forces that have the opposite sign as the applied strain. Most elements in a real structure are connected to finite stiffness, and so strain loading would cause both deformation and internal forces. Note that the effects of shear and bending strain loading are coupled.

For more information, see Topic “Internal Force and Stress Output” below, and also Chapter “Load Patterns” (page 321.)

Internal Force and Stress Output

The **Shell element internal forces** (also called **stress resultants**) are the forces and moments that result from integrating the stresses over the element thickness. For a homogeneous shell, these internal forces are:

- Membrane direct forces:

$$\begin{aligned}
 F_{11} & \int_{-\mathbf{th}/2}^{\mathbf{th}/2} \sigma_{11} dx_3 \\
 F_{22} & \int_{-\mathbf{th}/2}^{\mathbf{th}/2} \sigma_{22} dx_3
 \end{aligned}
 \tag{Eqns. 1}$$

- Membrane shear force:

$$F_{12} = \int_{-\mathbf{th}/2}^{\mathbf{th}/2} \tau_{12} dx_3$$

- Plate bending moments:

$$\begin{aligned}
 M_{11} & \int_{-\mathbf{thb}/2}^{\mathbf{thb}/2} x_3 \sigma_{11} dx_3 \\
 M_{22} & \int_{-\mathbf{thb}/2}^{\mathbf{thb}/2} x_3 \sigma_{22} dx_3
 \end{aligned}$$

- Plate twisting moment:

$$M_{12} = \int_{-\mathbf{thb}/2}^{\mathbf{thb}/2} x_3 \tau_{12} dx_3$$

- Plate transverse shear forces:

$$\begin{aligned}
 V_{13} & \int_{-\mathbf{thb}/2}^{\mathbf{thb}/2} \tau_{13} dx_3 \\
 V_{23} & \int_{-\mathbf{thb}/2}^{\mathbf{thb}/2} \tau_{23} dx_3
 \end{aligned}$$

where x_3 represents the thickness coordinate measured from the mid-surface of the element, \mathbf{th} is the membrane thickness, and \mathbf{thb} is the plate-bending thickness.

For a layered shell, the definitions are the same, except that the integrals of the stresses are now summed over all layers, and x_3 is always measured from the reference surface.

It is very important to note that these stress resultants are forces and moments *per unit of in-plane length*. They are present at every point on the mid-surface of the element.

For the thick-plate (Mindlin/Reissner) formulation of the homogeneous shell, and for the layered shell, the shear stresses are computed directly from the shearing deformation. For the thin-plate homogeneous shell, shearing deformation is assumed to be zero, so the transverse shear forces are computed instead from the moments using the equilibrium equations:

$$V_{13} = \frac{dM_{11}}{dx_1} - \frac{dM_{12}}{dx_2}$$

$$V_{23} \quad \frac{dM_{12}}{dx_1} \quad \frac{dM_{22}}{dx_2}$$

Where x_1 and x_2 are in-plane coordinates parallel to the local 1 and 2 axes.

The sign conventions for the stresses and internal forces are illustrated in Figure 51 (page 213). Stresses acting on a positive face are oriented in the positive direction of the element local coordinate axes. Stresses acting on a negative face are oriented in the negative direction of the element local coordinate axes.

A positive face is one whose outward normal (pointing away from element) is in the positive local 1 or 2 direction.

Positive internal forces correspond to a state of positive stress that is constant through the thickness. Positive internal moments correspond to a state of stress that varies linearly through the thickness and is positive at the bottom. Thus for a homogeneous shell:

$$\begin{array}{ll} 11 & \frac{F_{11}}{\mathbf{th}} \quad \frac{12 M_{11}}{\mathbf{thb}^3} x_3 \\ 22 & \frac{F_{22}}{\mathbf{th}} \quad \frac{12 M_{22}}{\mathbf{thb}^3} x_3 \\ 12 & \frac{F_{12}}{\mathbf{th}} \quad \frac{12 M_{12}}{\mathbf{thb}^3} x_3 \\ 13 & \frac{V_{13}}{\mathbf{thb}} \\ 23 & \frac{V_{23}}{\mathbf{thb}} \\ 33 & 0 \end{array} \quad (\text{Eqns. 2})$$

The transverse shear stresses given here are average values. The actual shear stress distribution is parabolic, being zero at the top and bottom surfaces and taking a maximum or minimum value at the mid-surface of the element.

The force and moment resultants are reported identically for homogeneous and layered shells. Stresses are reported for homogeneous shells at the top and bottom surfaces, and are linear in between. For the layered shell, stresses are reported in each layer at the integration points, and at the top, bottom, and center of the layer.

The stresses and internal forces are evaluated at the standard 2-by-2 Gauss integration points of the element and extrapolated to the joints. Although they are reported

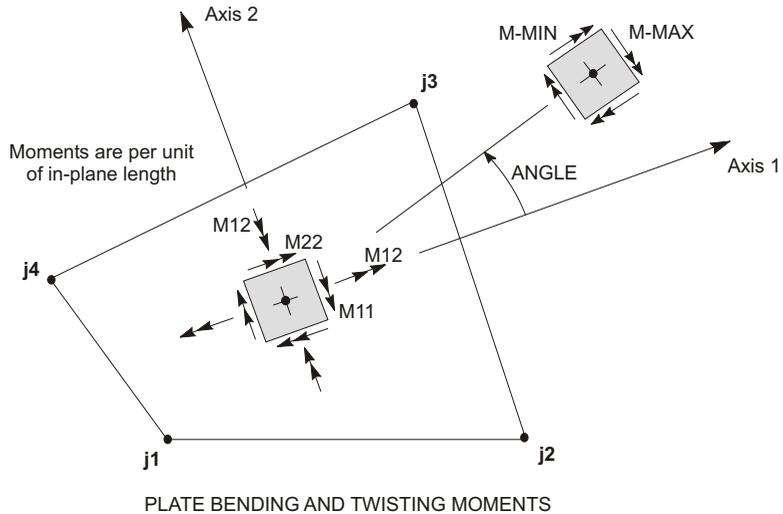
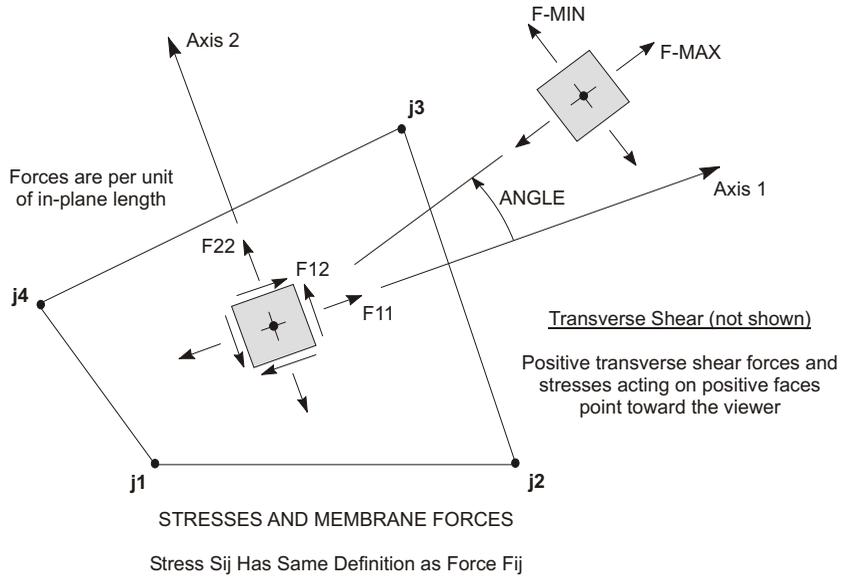


Figure 51
Shell Element Stresses and Internal Resultant Forces and Moments

at the joints, the stresses and internal forces exist over the whole element. See Cook, Malkus, and Plesha (1989) for more information.

Principal values and the associated principal directions are available for Load Cases and Load Combinations that are single valued. The angle given is measured counterclockwise (when viewed from the top) from the local 1 axis to the direction of the maximum principal value.

For more information, see Topic “Stresses and Strains” (page 71) in Chapter “Material Properties”.

The Plane Element

The **Plane** element is used to model plane-stress and plane-strain behavior in two-dimensional solids. The Plane element/object is one type of area object. Depending on the type of section properties you assign to an area, the object could also be used to model shell and axisymmetric solid behavior. These types of elements are discussed in the previous and following Chapters.

Advanced Topics

- Overview
- Joint Connectivity
- Degrees of Freedom
- Local Coordinate System
- Stresses and Strains
- Section Properties
- Mass
- Self-Weight Load
- Gravity Load
- Surface Pressure Load
- Pore Pressure Load

- Temperature Load
- Stress Output

Overview

The Plane element is a three- or four-node element for modeling two-dimensional solids of uniform thickness. It is based upon an isoparametric formulation that includes four optional incompatible bending modes. The element should be planar; if it is not, it is formulated for the projection of the element upon an average plane calculated for the element.

The incompatible bending modes significantly improve the bending behavior of the element if the element geometry is of a rectangular form. Improved behavior is exhibited even with non-rectangular geometry.

Structures that can be modeled with this element include:

- Thin, planar structures in a state of plane stress
- Long, prismatic structures in a state of plane strain

The stresses and strains are assumed not to vary in the thickness direction.

For plane-stress, the element has no out-of-plane stiffness. For plane-strain, the element can support loads with anti-plane shear stiffness.

Each Plane element has its own local coordinate system for defining Material properties and loads, and for interpreting output. Temperature-dependent, orthotropic material properties are allowed. Each element may be loaded by gravity (in any direction); surface pressure on the side faces; pore pressure within the element; and loads due to temperature change.

An 2×2 numerical integration scheme is used for the Plane. Stresses in the element local coordinate system are evaluated at the integration points and extrapolated to the joints of the element. An approximate error in the stresses can be estimated from the difference in values calculated from different elements attached to a common joint. This will give an indication of the accuracy of the finite element approximation and can then be used as the basis for the selection of a new and more accurate finite element mesh.

Joint Connectivity

The joint connectivity and face definition is identical for all area objects, i.e., the Shell, Plane, and Asolid elements. See Topic “Joint Connectivity” (page 180) in Chapter “The Shell Element” for more information.

The Plane element is intended to be planar. If you define a four-node element that is not planar, an average plane will be fit through the four joints, and the projection of the element onto this plane will be used.

Degrees of Freedom

The Plane element activates the three translational degrees of freedom at each of its connected joints. Rotational degrees of freedom are not activated.

The plane-stress element contributes stiffness only to the degrees of freedom in the plane of the element. It is necessary to provide restraints or other supports for the translational degrees of freedom that are normal to this plane; otherwise, the structure will be unstable.

The plane-strain element models anti-plane shear, i.e., shear that is normal to the plane of the element, in addition to the in-plane behavior. Thus stiffness is created for all three translational degrees of freedom.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom” for more information.

Local Coordinate System

The **element local coordinate system** is identical for all area objects, i.e., the Shell, Plane, and Asolid elements. See Topics “Local Coordinate System” (page 185) and “Advanced Local Coordinate System” (page 186) in Chapter “The Shell Element” for more information.

Stresses and Strains

The Plane element models the mid-plane of a structure having uniform thickness, and whose stresses and strains do not vary in the thickness direction.

Plane-stress is appropriate for structures that are *thin* compared to their planar dimensions. The thickness normal stress (σ_{33}) is assumed to be zero. The thickness normal strain (ϵ_{33}) may not be zero due to Poisson effects. Transverse shear stresses (τ_{12}, τ_{13}) and shear strains (γ_{12}, γ_{13}) are assumed to be zero. Displacements in the thickness (local 3) direction have no effect on the element.

Plane-strain is appropriate for structures that are *thick* compared to their planar dimensions. The thickness normal strain (ϵ_{33}) is assumed to be zero. The thickness normal stress (σ_{33}) may not be zero due to Poisson effects. Transverse shear stresses (τ_{12}, τ_{13}) and shear strains (γ_{12}, γ_{13}) are dependent upon displacements in the thickness (local 3) direction.

See Topic “Stresses and Strains” (page 71) in Chapter “Material Properties” for more information.

Section Properties

A **Plane Section** is a set of material and geometric properties that describe the cross-section of one or more Plane elements. Sections are defined independently of the Plane elements, and are assigned to the area objects.

Section Type

When defining an area section, you have a choice of three basic element types:

- Plane (stress or strain) – the subject of this Chapter, a two-dimensional solid, with translational degrees of freedom, capable of supporting forces but not moments.
- Shell – shell, plate, or membrane, with translational and rotational degrees of freedom, capable of supporting forces and moments. This element is covered in Chapter “The Shell Element” (page 177).
- Asolid – axisymmetric solid, with translational degrees of freedom, capable of supporting forces but not moments. This element is covered in Chapter “The Asolid Element” (page 225).

For Plane sections, you may choose one of the following sub-types of behavior:

- Plane stress
- Plane strain, including anti-plane shear

Material Properties

The material properties for each Plane element are specified by reference to a previously-defined Material. Orthotropic properties are used, even if the Material selected was defined as anisotropic. The material properties used by the Plane element are:

- The moduli of elasticity, **e1**, **e2**, and **e3**
- The shear modulus, **g12**
- For plane-strain only, the shear moduli, **g13** and **g23**
- The Poisson's ratios, **u12**, **u13** and **u23**
- The coefficients of thermal expansion, **a1**, **a2**, and **a3**
- The mass density, **m**, for computing element mass
- The weight density, **w**, for computing Self-Weight and Gravity Loads

The properties **e3**, **u13**, **u23**, and **a3** are not used for plane stress. They are used to compute the thickness-normal stress (σ_{33}) in plane strain.

All material properties (except the densities) are obtained at the material temperature of each individual element.

See Chapter “Material Properties” (page 69) for more information.

Material Angle

The material local coordinate system and the element (Plane Section) local coordinate system need not be the same. The local 3 directions always coincide for the two systems, but the material 1 axis and the element 1 axis may differ by the angle **a** as shown in Figure 52 (page 220). This angle has no effect for isotropic material properties since they are independent of orientation.

See Topic “Local Coordinate System” (page 70) in Chapter “Material Properties” for more information.

Thickness

Each Plane Section has a uniform thickness, **th**. This may be the actual thickness, particularly for plane-stress elements; or it may be a representative portion, such as a unit thickness of an infinitely-thick plane-strain element.

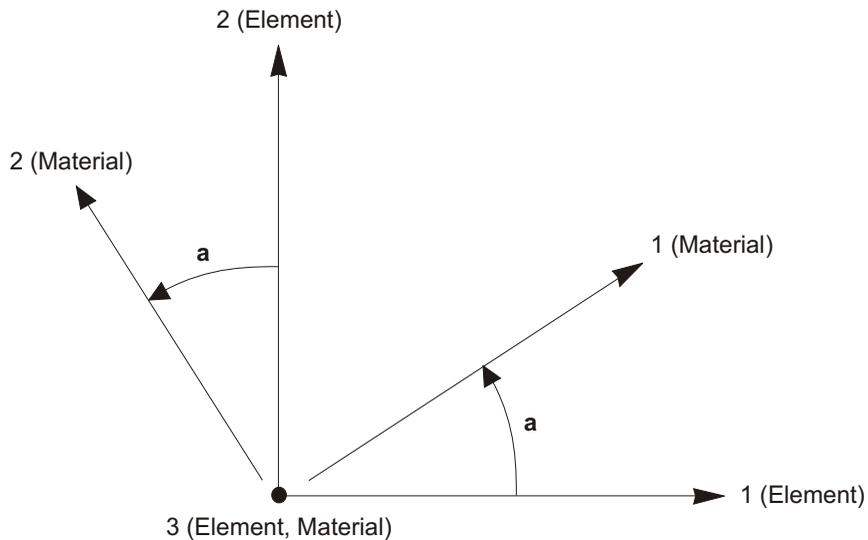


Figure 52
Plane Element Material Angle

The element thickness is used for calculating the element stiffness, mass, and loads. Hence, joint forces computed from the element are proportional to this thickness.

Incompatible Bending Modes

By default each Plane element includes four incompatible bending modes in its stiffness formulation. These incompatible bending modes significantly improve the bending behavior in the plane of the element if the element geometry is of a rectangular form. Improved behavior is exhibited even with non-rectangular geometry.

If an element is severely distorted, the inclusion of the incompatible modes should be suppressed. The element then uses the standard isoparametric formulation. Incompatible bending modes may also be suppressed in cases where bending is not important, such as in typical geotechnical problems.

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Plane element is lumped at the element joints. *No inertial effects are considered within the element itself.*

The total mass of the element is equal to the integral over the plane of the element of the mass density, \mathbf{m} , multiplied by the thickness, \mathbf{th} . The total mass is apportioned to the joints in a manner that is proportional to the diagonal terms of the consistent mass matrix. See Cook, Malkus, and Plesha (1989) for more information. The total mass is applied to each of the three translational degrees of freedom (UX, UY, and UZ) even when the element contributes stiffness to only two of these degrees of freedom.

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties.”
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Load Cases” (page 341).

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For a Plane element, the self-weight is a force that is uniformly distributed over the plane of the element. The magnitude of the self-weight is equal to the weight density, \mathbf{w} , multiplied by the thickness, \mathbf{th} .

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of \mathbf{w} .
- See Topic “Thickness” (page 219) in this Chapter for the definition of \mathbf{th} .
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Plane element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 221) in this Chapter for the definition of self-weight for the Plane element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Surface Pressure Load

The Surface Pressure Load is used to apply external pressure loads upon any of the three or four side faces of the Plane element. The definition of these faces is shown in Figure 42 (page 181). Surface pressure always acts normal to the face. Positive pressures are directed toward the interior of the element.

The pressure may be constant over a face or interpolated from values given at the joints. The values given at the joints are obtained from Joint Patterns, and need not be the same for the different faces. Joint Patterns can be used to easily apply hydrostatic pressures.

The pressure acting on a side is multiplied by the thickness, **th**, integrated along the length of the side, and apportioned to the two or three joints on that side.

See Chapter “Load Patterns” (page 321) for more information.

Pore Pressure Load

The Pore Pressure Load is used to model the drag and buoyancy effects of a fluid *within* a solid medium, such as the effect of water upon the solid skeleton of a soil.

Scalar fluid-pressure values are given at the element joints by Joint Patterns, and interpolated over the element. The total force acting on the element is the integral of the gradient of this pressure field over the plane of the element, multiplied by the thickness, **th**. This force is apportioned to each of the joints of the element. The forces are typically directed from regions of high pressure toward regions of low pressure.

See Chapter “Load Patterns” (page 321) for more information.

Temperature Load

The Temperature Load creates thermal strain in the Plane element. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the element. The temperature change is measured from the element Reference Temperature to the element Load Temperature. Temperature changes are assumed to be constant through the element thickness.

See Chapter “Load Patterns” (page 321) for more information.

Stress Output

The Plane element stresses are evaluated at the standard 2-by-2 Gauss integration points of the element and extrapolated to the joints. See Cook, Malkus, and Plesha (1989) for more information.

Principal values and their associated principal directions in the element local 1-2 plane are also computed for single-valued Load Cases. The angle given is measured counterclockwise (when viewed from the +3 direction) from the local 1 axis to the direction of the maximum principal value.

For more information:

- See Chapter “Load Patterns” (page 321).
- See Chapter “Load Cases” (page 341).

The Asolid Element

The **Asolid** element is used to model axisymmetric solids under axisymmetric loading.

Advanced Topics

- Overview
- Joint Connectivity
- Degrees of Freedom
- Local Coordinate System
- Stresses and Strains
- Section Properties
- Mass
- Self-Weight Load
- Gravity Load
- Surface Pressure Load
- Pore Pressure Load
- Temperature Load
- Rotate Load

- Stress Output

Overview

The Asolid element is a three- or four-node element for modeling axisymmetric structures under axisymmetric loading. It is based upon an isoparametric formulation that includes four optional incompatible bending modes.

The element models a representative two-dimensional cross section of the three-dimensional axisymmetric solid. The axis of symmetry may be located arbitrarily in the model. Each element should lie fully in a plane containing the axis of symmetry. If it does not, it is formulated for the projection of the element upon the plane containing the axis of symmetry and the center of the element.

The geometry, loading, displacements, stresses, and strains are assumed not to vary in the circumferential direction. Any displacements that occur in the circumferential direction are treated as axisymmetric torsion.

The use of incompatible bending modes significantly improves the in-plane bending behavior of the element if the element geometry is of a rectangular form. Improved behavior is exhibited even with non-rectangular geometry.

Each Asolid element has its own local coordinate system for defining Material properties and loads, and for interpreting output. Temperature-dependent, orthotropic material properties are allowed. Each element may be loaded by gravity (in any direction); centrifugal force; surface pressure on the side faces; pore pressure within the element; and loads due to temperature change.

An 2×2 numerical integration scheme is used for the Asolid. Stresses in the element local coordinate system are evaluated at the integration points and extrapolated to the joints of the element. An approximate error in the stresses can be estimated from the difference in values calculated from different elements attached to a common joint. This will give an indication of the accuracy of the finite element approximation and can then be used as the basis for the selection of a new and more accurate finite element mesh.

Joint Connectivity

The joint connectivity and face definition is identical for all area objects, i.e., the Shell, Plane, and Asolid elements. See Topic “Joint Connectivity” (page 180) in Chapter “The Shell Element” for more information.

The Asolid element is intended to be planar and to lie in a plane that contains the axis of symmetry. If not, a plane is found that contains the axis of symmetry and the center of the element, and the projection of the element onto this plane will be used.

Joints for a given element may not lie on opposite sides of the axis of symmetry. They may lie on the axis of symmetry and/or to one side of it.

Degrees of Freedom

The Asolid element activates the three translational degrees of freedom at each of its connected joints. Rotational degrees of freedom are not activated.

Stiffness is created for all three degrees of freedom. Degrees of freedom in the plane represent the radial and axial behavior. The normal translation represents circumferential torsion.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom” for more information.

Local Coordinate System

The **element local coordinate system** is identical for all area objects, i.e., the Shell, Plane, and Asolid elements. See Topics “Local Coordinate System” (page 185) and “Advanced Local Coordinate System” (page 186) in Chapter “The Shell Element” for more information.

The local 3 axis is normal to the plane of the element, and is the negative of the circumferential direction. The 1-2 plane is the same as the radial-axial plane, although the orientation of the local axes is not restricted to be parallel to the radial and axial axes.

The radial direction runs perpendicularly from the axis of symmetry to the center of the element. The axial direction is parallel to the axis of symmetry, with the positive sense being upward when looking along the circumferential (-3) direction with the radial direction pointing to the right.

Stresses and Strains

The **Asolid** element models the mid-plane of a representative sector of an axisymmetric structure whose stresses and strains do not vary in the circumferential direction.

Displacements in the local 1-2 plane cause in-plane strains ($\epsilon_{11}, \epsilon_{22}, \epsilon_{12}$) and stresses ($\sigma_{11}, \sigma_{22}, \sigma_{12}$).

Displacements in the radial direction also cause circumferential normal strains:

$$\epsilon_{33} = \frac{u_r}{r}$$

where u_r is the radial displacement, and r is the radius at the point in question. The circumferential normal stress (σ_{33}) is computed as usual from the three normal strains.

Displacements in the circumferential (local 3) direction cause only torsion, resulting in circumferential shear strains ($\epsilon_{12}, \epsilon_{13}$) and stresses (σ_{12}, σ_{13}).

See Topic “Stresses and Strains” (page 71) in Chapter “Material Properties” for more information.

Section Properties

An **Asolid Section** is a set of material and geometric properties that describe the cross-section of one or more **Asolid** elements. Sections are defined independently of the **Asolid** elements, and are assigned to the area objects.

Section Type

When defining an area section, you have a choice of three basic element types:

- **Asolid** – the subject of this Chapter, an axisymmetric solid, with translational degrees of freedom, capable of supporting forces but not moments.
- **Plane** (stress or strain) – a two-dimensional solid, with translational degrees of freedom, capable of supporting forces but not moments. This element is covered in Chapter “The Plane Element” (page 215).

- Shell – shell, plate, or membrane, with translational and rotational degrees of freedom, capable of supporting forces and moments. This element is covered in Chapter “The Shell Element” (page 177).

After selecting an Asolid type of section, you must supply the rest of the data described below.

Material Properties

The material properties for each Asolid element are specified by reference to a previously-defined Material. Orthotropic properties are used, even if the Material selected was defined as anisotropic. The material properties used by the Asolid element are:

- The moduli of elasticity, **e1**, **e2**, and **e3**
- The shear moduli, **g12**, **g13**, and **g23**
- The Poisson’s ratios, **u12**, **u13** and **u23**
- The coefficients of thermal expansion, **a1**, **a2**, and **a3**
- The mass density, **m**, for computing element mass
- The weight density, **w**, for computing Self-Weight and Gravity Loads

All material properties (except the densities) are obtained at the material temperature of each individual element.

See Chapter “Material Properties” (page 69) for more information.

Material Angle

The material local coordinate system and the element (Asolid Section) local coordinate system need not be the same. The local 3 directions always coincide for the two systems, but the material 1 axis and the element 1 axis may differ by the angle **a** as shown in Figure 53 (page 230). This angle has no effect for isotropic material properties since they are independent of orientation.

See Topic “Local Coordinate System” (page 70) in Chapter “Material Properties” for more information.

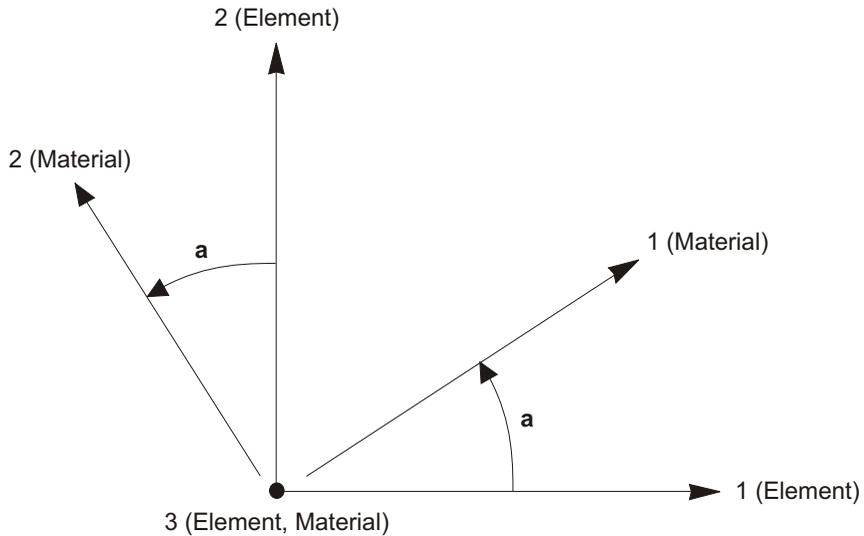


Figure 53
Asolid Element Material Angle

Axis of Symmetry

For each Asolid Section, you may select an axis of symmetry. This axis is specified as the Z axis of an alternate coordinate system that you have defined. All Asolid elements that use a given Asolid Section will have the same axis of symmetry.

For most modeling cases, you will only need a single axis of symmetry. However, if you want to have multiple axes of symmetry in your model, just set up as many alternate coordinate systems as needed for this purpose and define corresponding Asolid Section properties.

You should be aware that it is almost impossible to make a sensible model that connects Asolid elements with other element types, or that connects together Asolid elements using different axes of symmetry. The practical application of having multiple axes of symmetry is to have multiple *independent* axisymmetric structures in the same model.

See Topic “Alternate Coordinate Systems” (page 16) in Chapter “Coordinate Systems” for more information.

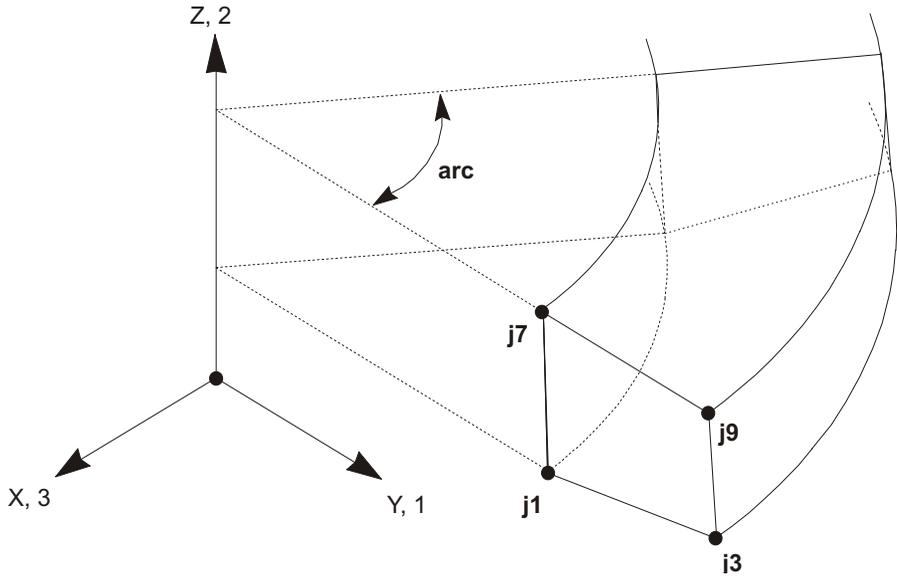


Figure 54
Asolid Element Local Coordinate System and Arc Definition

Arc and Thickness

The Asolid element represents a solid that is created by rotating the element's planar shape through 360° about the axis of symmetry. However, the analysis considers only a representative sector of the solid. You can specify the size of the sector, in degrees, using the parameter **arc**. For example, **arc**=360 models the full structure, and **arc**=90 models one quarter of it. See Figure 54 (page 231). Setting **arc**=0, the default, models a one-radian sector. One radian is the same as $180^\circ/\pi$, or approximately 57.3° .

The element "thickness" (circumferential extent), h , increases with the radial distance, r , from the axis of symmetry:

$$h = \frac{\text{arc}}{180} r$$

Clearly the thickness varies over the plane of the element.

The element thickness is used for calculating the element stiffness, mass, and loads. Hence, joint forces computed from the element are proportional to **arc**.

Incompatible Bending Modes

By default each Asolid element includes four incompatible bending modes in its stiffness formulation. These incompatible bending modes significantly improve the bending behavior in the plane of the element if the element geometry is of a rectangular form. Improved behavior is exhibited even with non-rectangular geometry.

If an element is severely distorted, the inclusion of the incompatible modes should be suppressed. The element then uses the standard isoparametric formulation. Incompatible bending modes may also be suppressed in cases where bending is not important, such as in typical geotechnical problems.

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Asolid element is lumped at the element joints. *No inertial effects are considered within the element itself.*

The total mass of the element is equal to the integral over the plane of the element of the product of the mass density, \mathbf{m} , multiplied by the thickness, h . The total mass is apportioned to the joints in a manner that is proportional to the diagonal terms of the consistent mass matrix. See Cook, Malkus, and Plesha (1989) for more information. The total mass is applied to each of the three translational degrees of freedom (UX, UY, and UZ).

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties.”
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Load Cases” (page 341).

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For an Asolid element, the self-weight is a force that is distributed over the plane of the element. The magnitude of the self-weight is equal to the weight density, \mathbf{w} , multiplied by the thickness, h .

Self-Weight Load always acts downward, in the global $-Z$ direction. If the downward direction corresponds to the radial or circumferential direction of an Asolid

element, the Self-Weight Load for that element will be zero, since self-weight acting in these directions is not axisymmetric. Non-zero Self-Weight Load will only exist for elements whose axial direction is vertical.

You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of w .
- See Subtopic “Arc and Thickness” (page 231) in this Chapter for the definition of h .
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Asolid element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element. However, only the components of Gravity load acting in the axial direction of an Asolid element will be non-zero. Components in the radial or circumferential direction will be set to zero, since gravity acting in these directions is not axisymmetric.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 232) in this Chapter for the definition of self-weight for the Asolid element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Surface Pressure Load

The Surface Pressure Load is used to apply external pressure loads upon any of the three or four side faces of the Asolid element. The definition of these faces is shown in Figure 42 (page 181). Surface pressure always acts normal to the face. Positive pressures are directed toward the interior of the element.

The pressure may be constant over a face or interpolated from values given at the joints. The values given at the joints are obtained from Joint Patterns, and need not be the same for the different faces. Joint Patterns can be used to easily apply hydrostatic pressures.

The pressure acting on a side is multiplied by the thickness, h , integrated along the length of the side, and apportioned to the two or three joints on that side.

See Chapter “Load Patterns” (page 321) for more information.

Pore Pressure Load

The Pore Pressure Load is used to model the drag and buoyancy effects of a fluid *within* a solid medium, such as the effect of water upon the solid skeleton of a soil.

Scalar fluid-pressure values are given at the element joints by Joint Patterns, and interpolated over the element. The total force acting on the element is the integral of the gradient of this pressure field, multiplied by the thickness h , over the plane of the element. This force is apportioned to each of the joints of the element. The forces are typically directed from regions of high pressure toward regions of low pressure.

See Chapter “Load Patterns” (page 321) for more information.

Temperature Load

The Temperature Load creates thermal strain in the A solid element. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the element. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis. Temperature changes are assumed to be constant through the element thickness.

See Chapter “Load Patterns” (page 321) for more information.

Rotate Load

Rotate Load is used to apply centrifugal force to A solid elements. Each element is assumed to rotate about its own axis of symmetry at a constant angular velocity.

The angular velocity creates a load on the element that is proportional to its mass, its distance from the axis of rotation, and the square of the angular velocity. This load acts in the positive radial direction, and is apportioned to each joint of the element. No Rotate Load will be produced by an element with zero mass density.

Since Rotate Loads assume a constant rate of rotation, it does not make sense to use a Load Pattern that contains Rotate Load in a time-history analysis unless that Load Pattern is applied quasi-statically (i.e., with a *very slow* time variation).

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties.”
- See Chapter “Load Patterns” (page 321).

Stress Output

The Asolid element stresses are evaluated at the standard 2-by-2 Gauss integration points of the element and extrapolated to the joints. See Cook, Malkus, and Plesha (1989) for more information.

Principal values and their associated principal directions in the element local 1-2 plane are also computed for single-valued Load Cases. The angle given is measured counterclockwise (when viewed from the +3 direction) from the local 1 axis to the direction of the maximum principal value.

For more information:

- See Chapter “Load Patterns” (page 321).
- See Chapter “Load Cases” (page 341).

The Solid Element

The **Solid** element is used to model three-dimensional solid structures.

Advanced Topics

- Overview
- Joint Connectivity
- Degrees of Freedom
- Local Coordinate System
- Advanced Local Coordinate System
- Stresses and Strains
- Solid Properties
- Mass
- Self-Weight Load
- Gravity Load
- Surface Pressure Load
- Pore Pressure Load
- Temperature Load
- Stress Output

Overview

The Solid element is an eight-node element for modeling three-dimensional structures and solids. It is based upon an isoparametric formulation that includes nine optional incompatible bending modes.

The incompatible bending modes significantly improve the bending behavior of the element if the element geometry is of a rectangular form. Improved behavior is exhibited even with non-rectangular geometry.

Each Solid element has its own local coordinate system for defining Material properties and loads, and for interpreting output. Temperature-dependent, anisotropic material properties are allowed. Each element may be loaded by gravity (in any direction); surface pressure on the faces; pore pressure within the element; and loads due to temperature change.

An $2 \times 2 \times 2$ numerical integration scheme is used for the Solid. Stresses in the element local coordinate system are evaluated at the integration points and extrapolated to the joints of the element. An approximate error in the stresses can be estimated from the difference in values calculated from different elements attached to a common joint. This will give an indication of the accuracy of the finite element approximation and can then be used as the basis for the selection of a new and more accurate finite element mesh.

Joint Connectivity

Each Solid element has six quadrilateral faces, with a joint located at each of the eight corners as shown in Figure 55 (page 239). It is important to note the relative position of the eight joints: the paths **j1-j2-j3** and **j5-j6-j7** should appear counter-clockwise when viewed along the direction from **j5** to **j1**. Mathematically stated, the three vectors:

- \mathbf{V}_{12} , from joints **j1** to **j2**,
- \mathbf{V}_{13} , from joints **j1** to **j3**,
- \mathbf{V}_{15} , from joints **j1** to **j5**,

must form a positive triple product, that is:

$$(\mathbf{V}_{12} \ \mathbf{V}_{13}) \cdot \mathbf{V}_{15} > 0$$

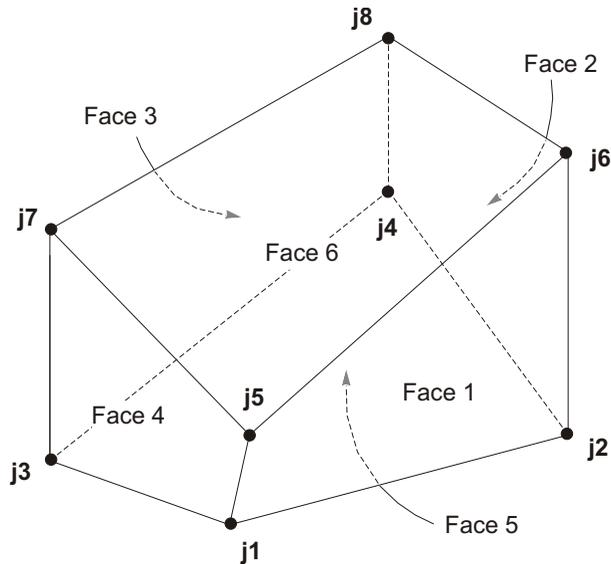


Figure 55
Solid Element Joint Connectivity and Face Definitions

The locations of the joints should be chosen to meet the following geometric conditions:

- The inside angle at each corner of the faces must be less than 180° . Best results will be obtained when these angles are near 90° , or at least in the range of 45° to 135° .
- The aspect ratio of an element should not be too large. This is the ratio of the longest dimension of the element to its shortest dimension. Best results are obtained for aspect ratios near unity, or at least less than four. The aspect ratio should not exceed ten.

These conditions can usually be met with adequate mesh refinement.

Degenerate Solids

Degenerate solids, such as wedges and tetrahedra, can be created by collapsing various sides of the element. This is done by specifying the same joint number for two or more of the eight corner nodes, so long as the ordering of the nodes remains the same. Examples include:

- Wedge (triangular bottom, triangular top): $j1, j2, j3 = j4, j5, j6, j7 = j8$
- Tetrahedron (triangular bottom, point top): $j1, j2, j3 = j4, j5 = j6 = j7 = j8$
- 7-node (rectangular bottom, triangular top): $j1, j2, j3, j4, j5, j6, j7 = j8$
- Pyramid (rectangular bottom, point top): $j1, j2, j3, j4, j5 = j6 = j7 = j8$

Other examples are possible, but these are the recommended configurations.

Degrees of Freedom

The Solid element activates the three translational degrees of freedom at each of its connected joints. Rotational degrees of freedom are not activated. This element contributes stiffness to all of these translational degrees of freedom.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom” for more information.

Local Coordinate System

Each Solid element has its own **element local coordinate system** used to define Material properties, loads and output. The axes of this local system are denoted 1, 2 and 3. By default these axes are identical to the global X, Y, and Z axes, respectively. Both systems are right-handed coordinate systems.

The default local coordinate system is adequate for most situations. However, for certain modeling purposes it may be useful to use element local coordinate systems that follow the geometry of the structure.

For more information:

- See Topic “Upward and Horizontal Directions” (page 13) in Chapter “Coordinate Systems.”
- See Topic “Advanced Local Coordinate System” (page 240) in this Chapter.

Advanced Local Coordinate System

By default, the element local 1-2-3 coordinate system is identical to the global X-Y-Z coordinate system, as described in the previous topic. In certain modeling situations it may be useful to have more control over the specification of the local coordinate system.

A variety of methods are available to define a solid-element local coordinate system. These may be used separately or together. Local coordinate axes may be defined to be parallel to arbitrary coordinate directions in an arbitrary coordinate system or to vectors between pairs of joints. In addition, the local coordinate system may be specified by a set of three element coordinate angles. These methods are described in the subtopics that follow.

For more information:

- See Chapter “Coordinate Systems” (page 11).
- See Topic “Local Coordinate System” (page 240) in this Chapter.

Reference Vectors

To define a solid-element local coordinate system you must specify two **reference vectors** that are parallel to one of the local coordinate planes. The **axis reference vector**, \mathbf{V}_a , must be parallel to one of the local axes ($I = 1, 2, \text{ or } 3$) in this plane and have a positive projection upon that axis. The **plane reference vector**, \mathbf{V}_p , must have a positive projection upon the other local axis ($j = 1, 2, \text{ or } 3, \text{ but } I \neq j$) in this plane, but need not be parallel to that axis. Having a positive projection means that the positive direction of the reference vector must make an angle of less than 90 with the positive direction of the local axis.

Together, the two reference vectors define a local axis, I , and a local plane, $i-j$. From this, the program can determine the third local axis, k , using vector algebra.

For example, you could choose the axis reference vector parallel to local axis 1 and the plane reference vector parallel to the local 1-2 plane ($I = 1, j = 2$). Alternatively, you could choose the axis reference vector parallel to local axis 3 and the plane reference vector parallel to the local 3-2 plane ($I = 3, j = 2$). You may choose the plane that is most convenient to define using the parameter **local**, which may take on the values 12, 13, 21, 23, 31, or 32. The two digits correspond to I and j , respectively. The default value is 31.

Defining the Axis Reference Vector

To define the axis reference vector, you must first specify or use the default values for:

- A coordinate direction **axdir** (the default is +Z)

- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system)

You may optionally specify:

- A pair of joints, **axveca** and **axvecb** (the default for each is zero, indicating the center of the element). If both are zero, this option is not used.

For each element, the axis reference vector is determined as follows:

1. A vector is found from joint **axveca** to joint **axvecb**. If this vector is of finite length, it is used as the reference vector \mathbf{V}_a
2. Otherwise, the coordinate direction **axdir** is evaluated at the center of the element in fixed coordinate system **csys**, and is used as the reference vector \mathbf{V}_a

Defining the Plane Reference Vector

To define the plane reference vector, you must first specify or use the default values for:

- A primary coordinate direction **pdirp** (the default is +X)
- A secondary coordinate direction **pdirs** (the default is +Y). Directions **pdirs** and **pdirp** should not be parallel to each other unless you are sure that they are not parallel to local axis 1
- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system). This will be the same coordinate system that was used to define the axis reference vector, as described above

You may optionally specify:

- A pair of joints, **plveca** and **plvecb** (the default for each is zero, indicating the center of the element). If both are zero, this option is not used.

For each element, the plane reference vector is determined as follows:

1. A vector is found from joint **plveca** to joint **plvecb**. If this vector is of finite length and is not parallel to local axis I , it is used as the reference vector \mathbf{V}_p
2. Otherwise, the primary coordinate direction **pdirp** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis I , it is used as the reference vector \mathbf{V}_p

3. Otherwise, the secondary coordinate direction **pldirs** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis I , it is used as the reference vector \mathbf{V}_p
4. Otherwise, the method fails and the analysis terminates. This will never happen if **pldirp** is not parallel to **pldirs**

A vector is considered to be parallel to local axis I if the sine of the angle between them is less than 10^{-3} .

Determining the Local Axes from the Reference Vectors

The program uses vector cross products to determine the local axes from the reference vectors. The three axes are represented by the three unit vectors \mathbf{V}_1 , \mathbf{V}_2 and \mathbf{V}_3 , respectively. The vectors satisfy the cross-product relationship:

$$\mathbf{V}_1 \times \mathbf{V}_2 = \mathbf{V}_3$$

The local axis \mathbf{V}_i is given by the vector \mathbf{V}_a after it has been normalized to unit length.

The remaining two axes, \mathbf{V}_j and \mathbf{V}_k , are defined as follows:

- If I and j permute in a positive sense, i.e., **local** = 12, 23, or 31, then:

$$\begin{aligned} \mathbf{V}_k &= \mathbf{V}_i \times \mathbf{V}_p \quad \text{and} \\ \mathbf{V}_j &= \mathbf{V}_k \times \mathbf{V}_i \end{aligned}$$

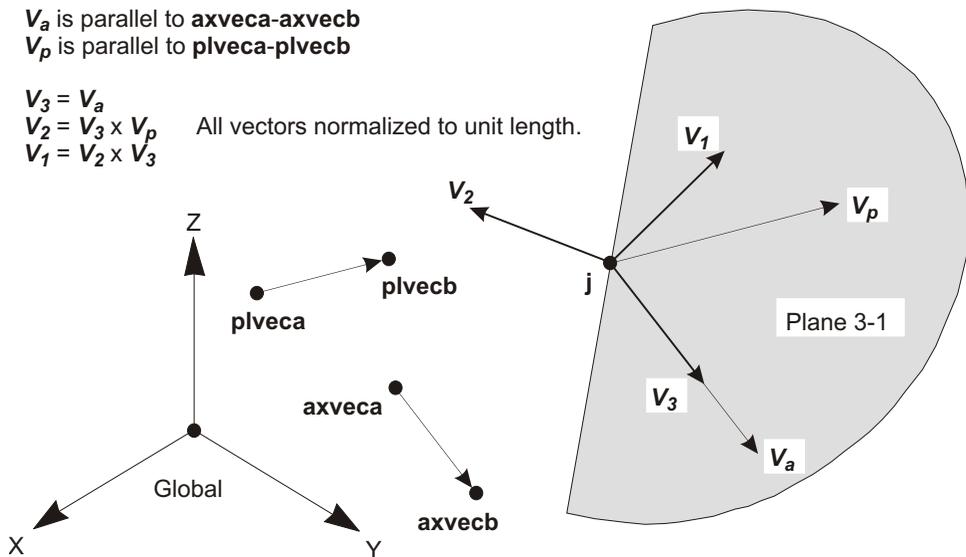
- If I and j permute in a negative sense, i.e., **local** = 21, 32, or 13, then:

$$\begin{aligned} \mathbf{V}_k &= \mathbf{V}_p \times \mathbf{V}_i \quad \text{and} \\ \mathbf{V}_j &= \mathbf{V}_i \times \mathbf{V}_k \end{aligned}$$

An example showing the determination of the element local coordinate system using reference vectors is given in Figure 56 (page 244).

Element Coordinate Angles

The solid-element local coordinate axes determined from the reference vectors may be further modified by the use of three **element coordinate angles**, denoted **a**, **b**, and **c**. In the case where the default reference vectors are used, the coordinate angles define the orientation of the element local coordinate system with respect to the global axes.

**Figure 56**

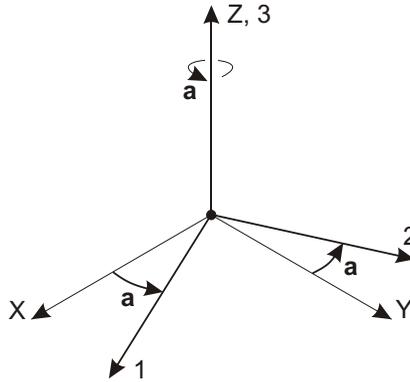
*Example of the Determination of the Solid Element Local Coordinate System Using Reference Vectors for **local=31**. Point **j** is the Center of the Element.*

The element coordinate angles specify rotations of the local coordinate system about its own current axes. The resulting orientation of the local coordinate system is obtained according to the following procedure:

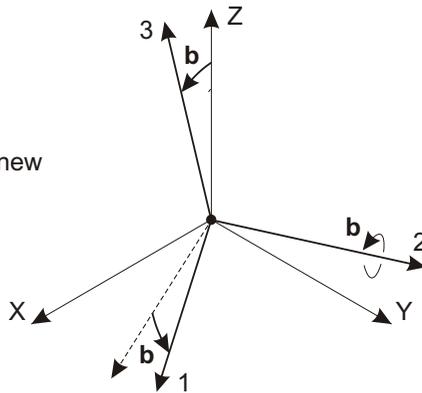
1. The local system is first rotated about its +3 axis by angle **a**
2. The local system is next rotated about its resulting +2 axis by angle **b**
3. The local system is lastly rotated about its resulting +1 axis by angle **c**

The order in which the rotations are performed is important. The use of coordinate angles to orient the element local coordinate system with respect to the global system is shown in Figure 4 (page 29).

Step 1: Rotation about local 3 axis by angle a



Step 2: Rotation about new local 2 axis by angle b



Step 3: Rotation about new local 1 axis by angle c

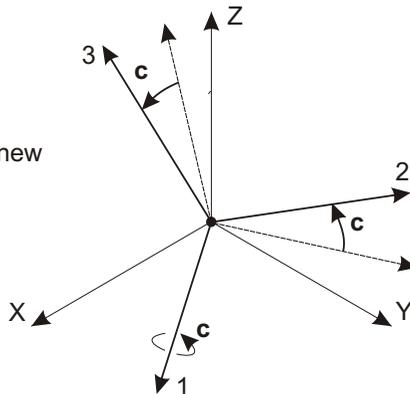


Figure 57
Use of Element Coordinate Angles to Orient the Solid Element Local Coordinate System

Stresses and Strains

The Solid element models a general state of stress and strain in a three-dimensional solid. All six stress and strain components are active for this element.

See Topic “Stresses and Strains” (page 71) in Chapter “Material Properties” for more information.

Solid Properties

A **Solid Property** is a set of material and geometric properties to be used by one or more Solid elements. Solid Properties are defined independently of the Solid elements/objects, and are assigned to the elements.

Material Properties

The material properties for each Solid Property are specified by reference to a previously-defined Material. Fully anisotropic material properties are used. The material properties used by the Solid element are:

- The moduli of elasticity, **e1**, **e2**, and **e3**
- The shear moduli, **g12**, **g13**, and **g23**
- All of the Poisson’s ratios, **u12**, **u13**, **u23**, ..., **u56**
- The coefficients of thermal expansion, **a1**, **a2**, **a3**, **a12**, **a13**, and **a23**
- The mass density, **m**, used for computing element mass
- The weight density, **w**, used for computing Self-Weight and Gravity Loads

All material properties (except the densities) are obtained at the material temperature of each individual element.

See Chapter “Material Properties” (page 69) for more information.

Material Angles

The material local coordinate system and the element (Property) local coordinate system need not be the same. The material coordinate system is oriented with respect to the element coordinate system using the three angles **a**, **b**, and **c** according to the following procedure:

- The material system is first aligned with the element system;

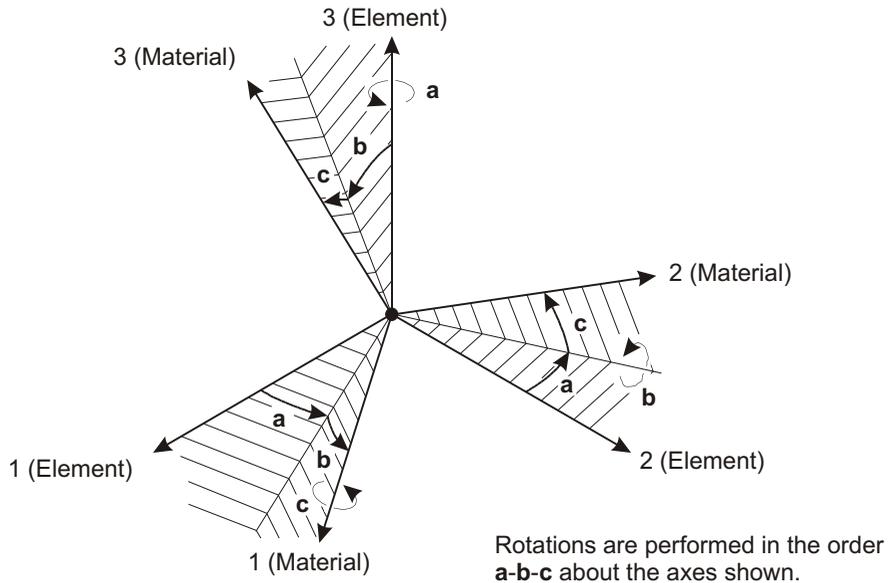


Figure 58
Solid Element Material Angles

- The material system is then rotated about its +3 axis by angle **a**;
- The material system is next rotated about the resulting +2 axis by angle **b**;
- The material system is lastly rotated about the resulting +1 axis by angle **c**;

This is shown in Figure 58 (page 247). These angles have no effect for isotropic material properties since they are independent of orientation.

See Topic “Local Coordinate System” (page 70) in Chapter “Material Properties” for more information.

Incompatible Bending Modes

By default each Solid element includes nine incompatible bending modes in its stiffness formulation. These incompatible bending modes significantly improve the bending behavior of the element if the element geometry is of a rectangular form. Improved behavior is exhibited even with non-rectangular geometry.

If an element is severely distorted, the inclusion of the incompatible modes should be suppressed. The element then uses the standard isoparametric formulation. In-

compatible bending modes may also be suppressed in cases where bending is not important, such as in typical geotechnical problems.

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Solid element is lumped at the element joints. *No inertial effects are considered within the element itself.*

The total mass of the element is equal to the integral of the mass density, \mathbf{m} , over the volume of the element. The total mass is apportioned to the joints in a manner that is proportional to the diagonal terms of the consistent mass matrix. See Cook, Malkus, and Plesha (1989) for more information. The total mass is applied to each of the three translational degrees of freedom (UX, UY, and UZ).

For more information:

- See Topic “Mass Density” (page 77) in Chapter “Material Properties.”
- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Load Cases” (page 341).

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For a Solid element, the self-weight is a force that is uniformly distributed over the volume of the element. The magnitude of the self-weight is equal to the weight density, \mathbf{w} .

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of \mathbf{w} .
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Solid element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 248) in this Chapter for the definition of self-weight for the Solid element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Surface Pressure Load

The Surface Pressure Load is used to apply external pressure loads upon any of the six faces of the Solid element. The definition of these faces is shown in Figure 55 (page 239). Surface pressure always acts normal to the face. Positive pressures are directed toward the interior of the element.

The pressure may be constant over a face or interpolated from values given at the joints. The values given at the joints are obtained from Joint Patterns, and need not be the same for the different faces. Joint Patterns can be used to easily apply hydrostatic pressures.

The pressure acting on a given face is integrated over the area of that face, and the resulting force is apportioned to the four corner joints of the face.

See Chapter “Load Patterns” (page 321) for more information.

Pore Pressure Load

The Pore Pressure Load is used to model the drag and buoyancy effects of a fluid *within* a solid medium, such as the effect of water upon the solid skeleton of a soil.

Scalar fluid-pressure values are given at the element joints by Joint Patterns, and interpolated over the element. The total force acting on the element is the integral of the gradient of this pressure field over the volume of the element. This force is ap-

portioned to each of the joints of the element. The forces are typically directed from regions of high pressure toward regions of low pressure.

See Chapter “Load Patterns” (page 321) for more information.

Temperature Load

The Temperature Load creates thermal strain in the Solid element. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the element. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis.

See Chapter “Load Patterns” (page 321) for more information.

Stress Output

The Solid element stresses are evaluated at the standard 2 x 2 x 2 Gauss integration points of the element and extrapolated to the joints. See Cook, Malkus, and Plesha (1989) for more information.

Principal values and their associated principal directions in the element local coordinate system are also computed for single-valued Load Cases and Load Combinations. Three direction cosines each are given for the directions of the maximum and minimum principal stresses. The direction of the middle principal stress is perpendicular to the maximum and minimum principal directions.

For more information:

- See Chapter “Load Patterns” (page 321).
- See Chapter “Load Cases” (page 341).

The Link/Support Element—Basic

The **Link** element is used to connect two joints together. The Support element is used to connect one joint to ground. Both element types use the same types of properties. Each Link or Support element may exhibit up to three different types of behavior: linear, nonlinear, and frequency-dependent, according to the types of properties assigned to that element and the type of analysis being performed.

This Chapter describes the basic and general features of the Link and Support elements and their linear behavior. The next Chapter describes advanced behavior, which can be nonlinear or frequency-dependent.

Advanced Topics

- Overview
- Joint Connectivity
- Zero-Length Elements
- Degrees of Freedom
- Local Coordinate System
- Advanced Local Coordinate System
- Internal Deformations
- Link/Support Properties

- Coupled Linear Property
- Mass
- Self-Weight Load
- Gravity Load
- Internal Force and Deformation Output

Overview

A Link element is a two-joint connecting link. A Support element is a one-joint grounded spring. Properties for both types of element are defined in the same way. Each element is assumed to be composed of six separate “springs,” one for each of six deformational degrees-of freedom (axial, shear, torsion, and pure bending).

There are two categories of Link/Support properties that can be defined: Linear/Nonlinear, and Frequency-Dependent. A Linear/Nonlinear property set must be assigned to each Link or Support element. The assignment of a Frequency-Dependent property set to a Link or Support element is optional.

All Linear/Nonlinear property sets contain linear properties that are used by the element for linear analyses, and for other types of analyses if no other properties are defined. Linear/Nonlinear property sets may have nonlinear properties that will be used for all nonlinear analyses, and for linear analyses that continue from nonlinear analyses.

Frequency-dependent property sets contain impedance (stiffness and damping) properties that will be used for all frequency-dependent analyses. If a Frequency-Dependent property has not been assigned to a Link/Support element, the linear properties for that element will be used for frequency-dependent analyses.

The types of nonlinear behavior that can be modeled with this element include:

- Viscoelastic damping
- Gap (compression only) and hook (tension only)
- Multi-linear uniaxial elasticity
- Uniaxial plasticity (Wen model)
- Multi-linear uniaxial plasticity with several types of hysteretic behavior: kinematic, Takeda, and pivot
- Biaxial-plasticity base isolator

- Friction-pendulum base isolator, with or without uplift prevention. This can also be used for modeling gap-friction contact behavior

Each element has its own local coordinate system for defining the force-deformation properties and for interpreting output.

Each Link/Support element may be loaded by gravity (in any direction).

Available output includes the deformation across the element, and the internal forces at the joints of the element.

Joint Connectivity

Each Link/Support element may take one of the following two configurations:

- A Link connecting two joints, **I** and **j**; it is permissible for the two joints to share the same location in space creating a zero-length element
- A Support connecting a single joint, **j**, to ground

Conversion from One-Joint Objects to Two-Joint Elements

During analysis, all one-joint Link/Support objects used for modeling are actually converted to two-joint Link/Support elements of zero length. The original joint, **j**, remains connected to the structure. A new joint, **I**, is generated at the same location and is fully restrained. Reactions are calculated at generated joint **I** but are reported at the original joint **j**. Conversely, ground displacement loads applied at joint **j** are transferred to joint **I**.

For the remainder of this chapter and the next, we will continue to refer to one-joint elements for convenience, and to clarify how the one-joint modeling objects behave.

Zero-Length Elements

The following types of Link/Support elements are considered to be of zero **length**:

- Single-joint Support elements
- Two-joint Link elements with the distance from joint **I** to joint **j** being less than or equal to the zero-length tolerance that you specify.

The length tolerance is set using the Auto Merge Tolerance in the graphical user interface. Two-joint elements having a length greater than the Auto Merge Tolerance are considered to be of finite length. Whether an element is of zero length or finite length affects the definition of the element local coordinate system, and the internal moments due to shear forces.

Degrees of Freedom

The Link/Support element always activates all six degrees of freedom at each of its one or two connected joints. To which joint degrees of freedom the element contributes stiffness depends upon the properties you assign to the element. You must ensure that restraints or other supports are provided to those joint degrees of freedom that receive no stiffness.

For more information:

- See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom.”
- See Topic “Link/Support Properties” (page 263) in this Chapter.

Local Coordinate System

Each Link/Support element has its own **element local coordinate system** used to define force-deformation properties and output. The axes of this local system are denoted 1, 2 and 3. The first axis is directed along the length of the element and corresponds to extensional deformation. The remaining two axes lie in the plane perpendicular to the element and have an orientation that you specify; these directions correspond to shear deformation.

It is important that you clearly understand the definition of the element local 1-2-3 coordinate system and its relationship to the global X-Y-Z coordinate system. Both systems are right-handed coordinate systems. It is up to you to define local systems which simplify data input and interpretation of results.

In most structures the definition of the element local coordinate system is extremely simple. The methods provided, however, provide sufficient power and flexibility to describe the orientation of Link/Support elements in the most complicated situations.

The simplest method, using the **default orientation** and the **Link/Support element coordinate angle**, is described in this topic. Additional methods for defining the Link/Support element local coordinate system are described in the next topic.

For more information:

- See Chapter “Coordinate Systems” (page 11) for a description of the concepts and terminology used in this topic.
- See Topic “Advanced Local Coordinate System” (page 256) in this Chapter.

Longitudinal Axis I

Local axis 1 is the longitudinal axis of the element, corresponding to extensional deformation. This axis is determined as follows:

- For elements of finite length this axis is automatically defined as the direction from joint **I** to joint **j**
- For zero-length elements the local 1 axis defaults to the +Z global coordinate direction (upward)

For the definition of zero-length elements, see Topic “Zero-Length Elements” (page 253) in this Chapter.

Default Orientation

The default orientation of the local 2 and 3 axes is determined by the relationship between the local 1 axis and the global Z axis. The procedure used here is identical to that for the Frame element:

- The local 1-2 plane is taken to be vertical, i.e., parallel to the Z axis
- The local 2 axis is taken to have an upward (+Z) sense unless the element is vertical, in which case the local 2 axis is taken to be horizontal along the global +X direction
- The local 3 axis is always horizontal, i.e., it lies in the X-Y plane

An element is considered to be vertical if the sine of the angle between the local 1 axis and the Z axis is less than 10^{-3} .

The local 2 axis makes the same angle with the vertical axis as the local 1 axis makes with the horizontal plane. This means that the local 2 axis points vertically upward for horizontal elements.

Coordinate Angle

The Link/Support element coordinate angle, **ang**, is used to define element orientations that are different from the default orientation. It is the angle through which the local 2 and 3 axes are rotated about the positive local 1 axis from the default orientation. The rotation for a positive value of **ang** appears counterclockwise when the local +1 axis is pointing toward you. The procedure used here is identical to that for the Frame element.

For vertical elements, **ang** is the angle between the local 2 axis and the horizontal +X axis. Otherwise, **ang** is the angle between the local 2 axis and the vertical plane containing the local 1 axis. See Figure 59 (page 257) for examples.

Advanced Local Coordinate System

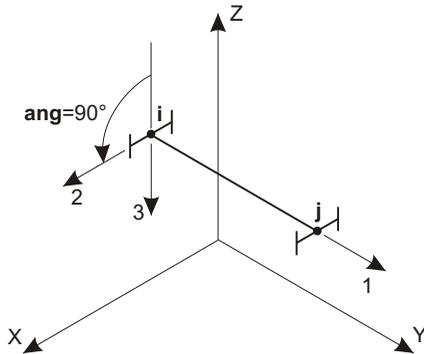
By default, the element local coordinate system is defined using the element coordinate angle measured with respect to the global +Z and +X directions, as described in the previous topic. In certain modeling situations it may be useful to have more control over the specification of the local coordinate system.

This topic describes how to define the orientation of the transverse local 2 and 3 axes with respect to an arbitrary reference vector when the element coordinate angle, **ang**, is zero. If **ang** is different from zero, it is the angle through which the local 2 and 3 axes are rotated about the positive local 1 axis from the orientation determined by the reference vector.

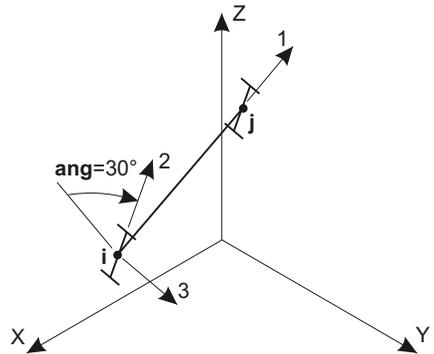
This topic also describes how to change the orientation of the local 1 axis from the default global +Z direction for zero-length elements. The local 1 axis is *always* directed from joint **I** to joint **j** for elements of *finite length*.

For more information:

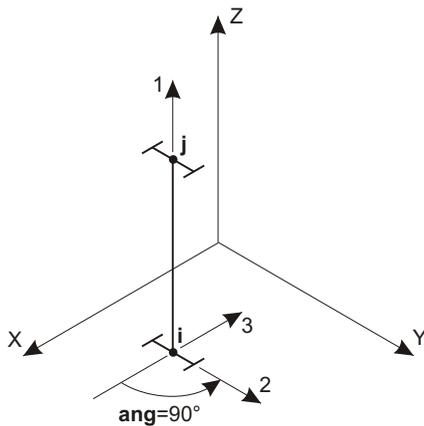
- See Chapter “Coordinate Systems” (page 11) for a description of the concepts and terminology used in this topic.
- See Topic “Local Coordinate System” (page 253) in this Chapter.



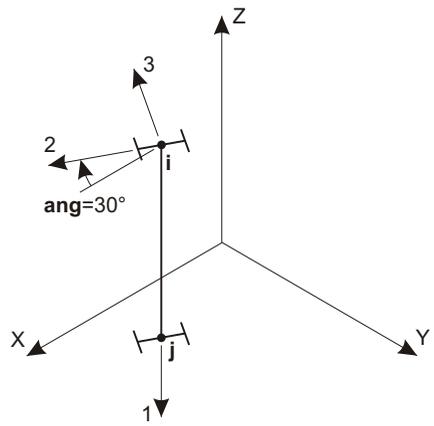
Local 1 Axis is Parallel to +Y Axis
Local 2 Axis is Rotated 90° from Z-1 Plane



Local 1 Axis is Not Parallel to X, Y, or Z Axes
Local 2 Axis is Rotated 30° from Z-1 Plane



Local 1 Axis is Parallel to +Z Axis
Local 2 Axis is Rotated 90° from X-1 Plane



Local 1 Axis is Parallel to -Z Axis
Local 2 Axis is Rotated 30° from X-1 Plane

Figure 59

The Link/Support Element Coordinate Angle with Respect to the Default Orientation

Axis Reference Vector

To define the local 1 axis for *zero-length* elements, you specify an **axis reference vector** that is parallel to and has the same positive sense as the desired local 1 axis. The axis reference vector has no effect upon finite-length elements.

To define the axis reference vector, you must first specify or use the default values for:

- A coordinate direction **axdir** (the default is +Z)
- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system). This will be the same coordinate system that is used to define the plane reference vector, as described below

You may optionally specify:

- A pair of joints, **axveca** and **axvecb** (the default for each is zero, indicating the center of the element). If both are zero, this option is not used

For each element, the axis reference vector is determined as follows:

1. A vector is found from joint **axveca** to joint **axvecb**. If this vector is of finite length, it is used as the reference vector \mathbf{V}_a
2. Otherwise, the coordinate direction **axdir** is evaluated at the center of the element in fixed coordinate system **csys**, and is used as the reference vector \mathbf{V}_a

The center of a zero-length element is taken to be at joint **j**.

The local 1 axis is given by the vector \mathbf{V}_a after it has been normalized to unit length.

Plane Reference Vector

To define the transverse local axes 2 and 3, you specify a **plane reference vector** that is parallel to the desired 1-2 or 1-3 plane. The procedure used here is identical to that for the Frame element.

The reference vector must have a positive projection upon the corresponding transverse local axis (2 or 3, respectively). This means that the positive direction of the reference vector must make an angle of less than 90° with the positive direction of the desired transverse axis.

To define the reference vector, you must first specify or use the default values for:

- A primary coordinate direction **pdirp** (the default is +Z)
- A secondary coordinate direction **pdirs** (the default is +X). Directions **pdirs** and **pdirp** should not be parallel to each other unless you are sure that they are not parallel to local axis 1

- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system). This will be the same coordinate system that was used to define the axis reference vector, as described above
- The local plane, **local**, to be determined by the reference vector (the default is 12, indicating plane 1-2)

You may optionally specify:

- A pair of joints, **plveca** and **plvecb** (the default for each is zero, indicating the center of the element). If both are zero, this option is not used

For each element, the reference vector is determined as follows:

1. A vector is found from joint **plveca** to joint **plvecb**. If this vector is of finite length and is not parallel to local axis 1, it is used as the reference vector V_p .
2. Otherwise, the primary coordinate direction **pldirp** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis 1, it is used as the reference vector V_p .
3. Otherwise, the secondary coordinate direction **pldirs** is evaluated at the center of the element in fixed coordinate system **csys**. If this direction is not parallel to local axis 1, it is used as the reference vector V_p .
4. Otherwise, the method fails and the analysis terminates. This will never happen if **pldirp** is not parallel to **pldirs**

A vector is considered to be parallel to local axis 1 if the sine of the angle between them is less than 10^{-3} .

The use of the Link/Support element coordinate angle in conjunction with coordinate directions that define the reference vector is illustrated in Figure 60 (page 260). The use of joints to define the reference vector is shown in Figure 61 (page 261).

Determining Transverse Axes 2 and 3

The program uses vector cross products to determine the transverse axes 2 and 3 once the reference vector has been specified. The three axes are represented by the three unit vectors V_1 , V_2 and V_3 , respectively. The vectors satisfy the cross-product relationship:

$$V_1 \times V_2 = V_3$$

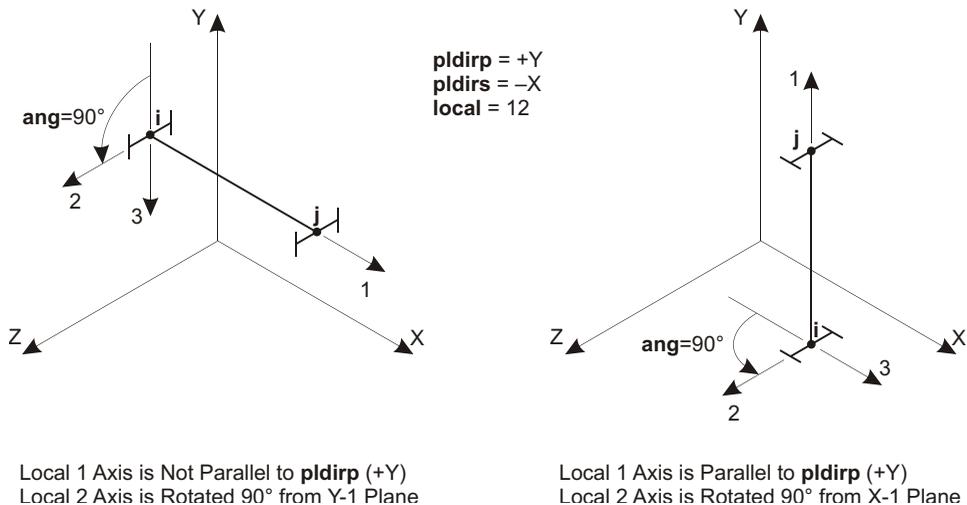


Figure 60
The Link/Support Element Coordinate Angle with Respect to Coordinate Directions

The transverse axes 2 and 3 are defined as follows:

- If the reference vector is parallel to the 1-2 plane, then:

$$\begin{matrix} \mathbf{V}_3 & \mathbf{V}_1 & \mathbf{V}_p & \text{and} \\ \mathbf{V}_2 & \mathbf{V}_3 & \mathbf{V}_1 \end{matrix}$$

- If the reference vector is parallel to the 1-3 plane, then:

$$\begin{matrix} \mathbf{V}_2 & \mathbf{V}_p & \mathbf{V}_1 & \text{and} \\ \mathbf{V}_3 & \mathbf{V}_1 & \mathbf{V}_2 \end{matrix}$$

In the common case where the reference vector is perpendicular to axis \mathbf{V}_1 , the transverse axis in the selected plane will be equal to \mathbf{V}_p .

Internal Deformations

Six independent **internal deformations** are defined for the Link/Support element. These are calculated from the relative displacements of joint **j** with respect to:

- Joint **I** for a two-joint element
- The ground for a single-joint element

The following two specifications are equivalent:

- (a) **local**=12, **plveca**=0, **plvecb**=100
- (b) **local**=13, **plveca**=101, **plvecb**=102

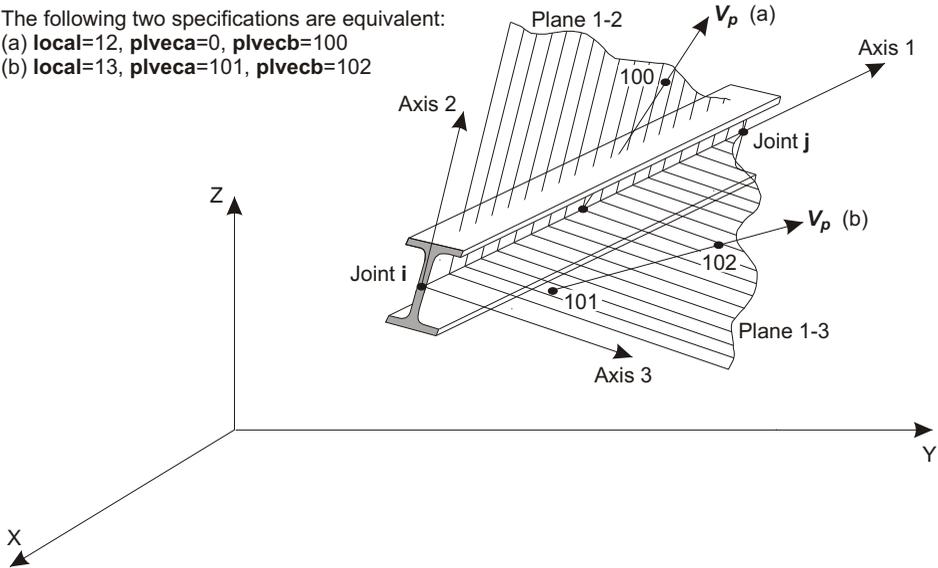


Figure 61

Using Joints to Define the Link/Support Element Local Coordinate System

For *two-joint Link/Support* elements the internal deformations are defined as:

- Axial: $d_{u1} = u_{1j} - u_{1i}$
- Shear in the 1-2 plane: $d_{u2} = u_{2j} - u_{2i} - \mathbf{d}j2 r_{3j} - (L - \mathbf{d}j2) r_{3i}$
- Shear in the 1-3 plane: $d_{u3} = u_{3j} - u_{3i} + \mathbf{d}j3 r_{2j} + (L - \mathbf{d}j3) r_{2i}$
- Torsion: $d_{r1} = r_{1j} - r_{1i}$
- Pure bending in the 1-3 plane: $d_{r2} = r_{2i} - r_{2j}$
- Pure bending in the 1-2 plane: $d_{r3} = r_{3j} - r_{3i}$

where:

- $u_{1i}, u_{2i}, u_{3i}, r_{1i}, r_{2i},$ and r_{3i} are the translations and rotations at joint **I**
- $u_{1j}, u_{2j}, u_{3j}, r_{1j}, r_{2j},$ and r_{3j} are the translations and rotations at joint **J**
- **dj2** is the distance you specify from joint **J** to the location where the shear deformation d_{u2} is measured (the default is zero, meaning at joint **J**)
- **dj3** is the distance you specify from joint **J** to the location where the shear deformation d_{u3} is measured (the default is zero, meaning at joint **J**)
- L is the length of the element

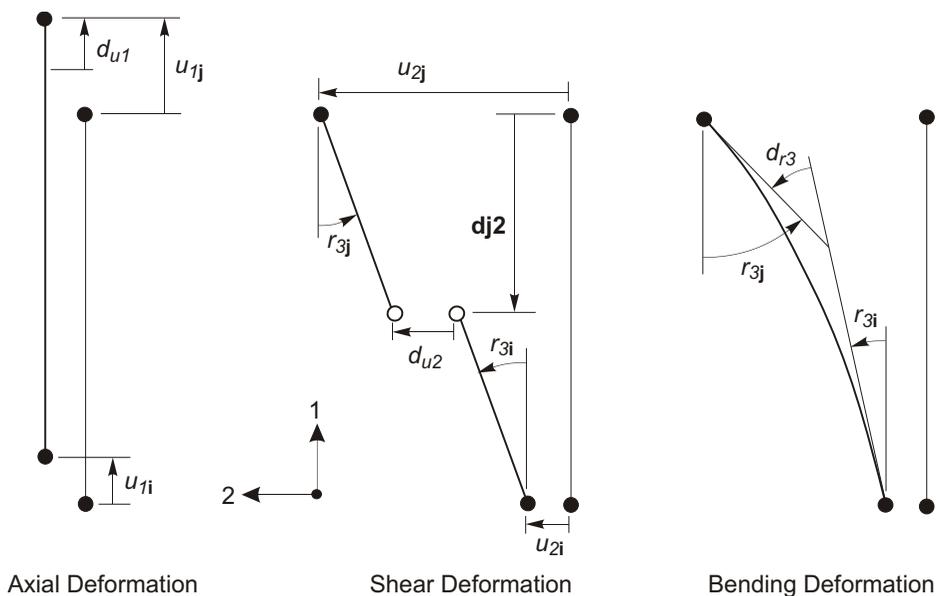


Figure 62
Internal Deformations for a Two-Joint Link Element

All translations, rotations, and deformations are expressed in terms of the element local coordinate system.

Note that shear deformation can be caused by rotations as well as translations. These definitions ensure that all deformations will be zero under rigid-body motions of the element.

Important! Note that **dj2** is the location where pure bending behavior is measured in the 1-2 plane, in other words, it is where the moment due to shear is taken to be zero. Likewise, **dj3** is the **location** where pure bending behavior is measured in the 1-3 plane.

It is important to note that the negatives of the rotations r_{2i} and r_{2j} have been used for the definition of shear and bending deformations in the 1-3 plane. This provides consistent definitions for shear and moment in both the Link/Support and Frame elements.

Three of these internal deformations are illustrated in Figure 62 (page 262).

For *one-joint grounded-spring* elements the internal deformations are the same as above, except that the translations and rotations at joint **I** are taken to be zero:

- Axial: $d_{u1} = u_{1j}$
- Shear in the 1-2 plane: $d_{u2} = u_{2j} - \mathbf{dj2} r_{3j}$
- Shear in the 1-3 plane: $d_{u3} = u_{3j} + \mathbf{dj3} r_{2j}$
- Torsion: $d_{r1} = r_{1j}$
- Pure bending in the 1-3 plane: $d_{r2} = -r_{2j}$
- Pure bending in the 1-2 plane: $d_{r3} = r_{3j}$

Link/Support Properties

A Link/Support Property is a set of structural properties that can be used to define the behavior of one or more Link or Support elements. Each Link/Support Property specifies the force-deformation relationships for the six internal deformations. Mass and weight properties may also be specified.

Link/Support Properties are defined independently of the Link and Support elements and are referenced during the definition of the elements.

There are two categories of Link/Support properties that can be defined:

- **Linear/Nonlinear.** A Linear/Nonlinear property set must be assigned to each Link or Support element.
- **Frequency-Dependent.** The assignment of a Frequency-Dependent property set to a Link or Support element is optional.

All Linear/Nonlinear property sets contain linear properties that are used by the element for linear analyses, and for other types of analyses if no other properties are defined. Linear/Nonlinear property sets may also have nonlinear properties that will be used for all nonlinear analyses, and for linear analyses that continue from nonlinear analyses.

Frequency-dependent property sets contain impedance (stiffness and damping) properties that will be used for all frequency-dependent analyses. If a Frequency-Dependent property has not been assigned to a Link/Support element, the linear properties for that element will be used for frequency-dependent analyses.

This is summarized in the table of Figure 63 (page 264).

Load Case Type	Load Case Initial Conditions	Element has Nonlinear Properties?	Element has Freq. Dep. Properties?	Actual Property Used
Linear	Zero	Yes or No	Yes or No	Linear
	Nonlinear Case	No	Yes or No	Linear
		Yes	Yes or No	Nonlinear
Nonlinear	Any	No	Yes or No	Linear
		Yes	Yes or No	Nonlinear
Frequency Dependent	Zero	Yes or No	No	Linear
			Yes	Freq. Dep.
	Nonlinear Case	No	No	Linear
		Yes	No	Nonlinear
		Yes or No	Yes	Freq. Dep.

Figure 63
Link/Support Stiffness Properties Actually Used for Different Types of Analysis

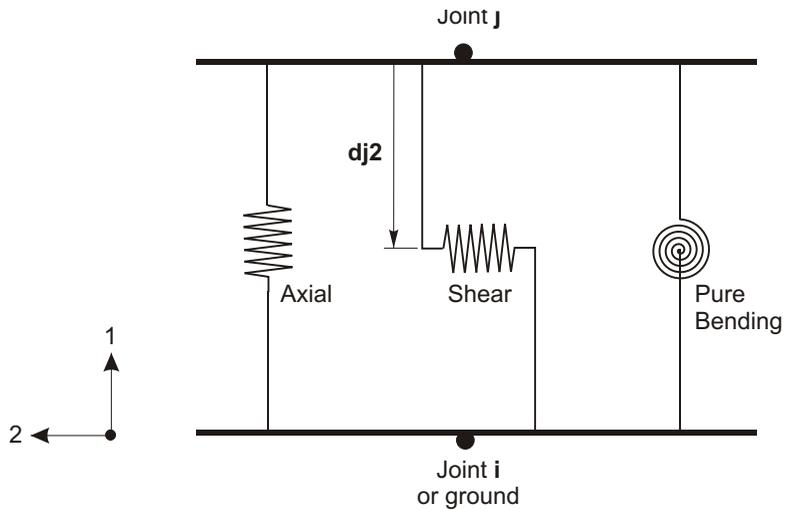
Local Coordinate System

Link/Support Properties are defined with respect to the local coordinate system of the Link or Support element. The local 1 axis is the longitudinal direction of the element and corresponds to extensional and torsional deformations. The local 2 and 3 directions correspond to shear and bending deformations.

See Topic “Local Coordinate System” (page 253) in this Chapter.

Internal Spring Hinges

Each Link/Support Property is assumed to be composed of six internal “springs” or “Hinges,” one for each of six internal deformations. Each “spring” may actually

**Figure 64**

Three of the Six Independent Spring Hinges in a Link/Support Element

consist of several components, including springs and dashpots. The force-deformation relationships of these springs may be coupled or independent of each other.

Figure 64 (page 265) shows the springs for three of the deformations: axial, shear in the 1-2 plane, and pure-bending in the 1-2 plane. It is important to note that the shear spring is located a distance d_{j2} from joint **j**. All shear deformation is assumed to occur in this spring; the links connecting this spring to the joints (or ground) are rigid in shear. Deformation of the shear spring can be caused by rotations as well as translations at the joints. The force in this spring will produce a linearly-varying moment along the length. This moment is taken to be zero at the shear spring, which acts as a moment hinge. The moment due to shear is independent of, and additive to, the constant moment in the element due to the pure-bending spring.

The other three springs that are not shown are for torsion, shear in the 1-3 plane, and pure-bending in the 1-3 plane. The shear spring is located a distance d_{j3} from joint **j**.

The values of d_{j2} and d_{j3} may be different, although normally they would be the same for most elements.

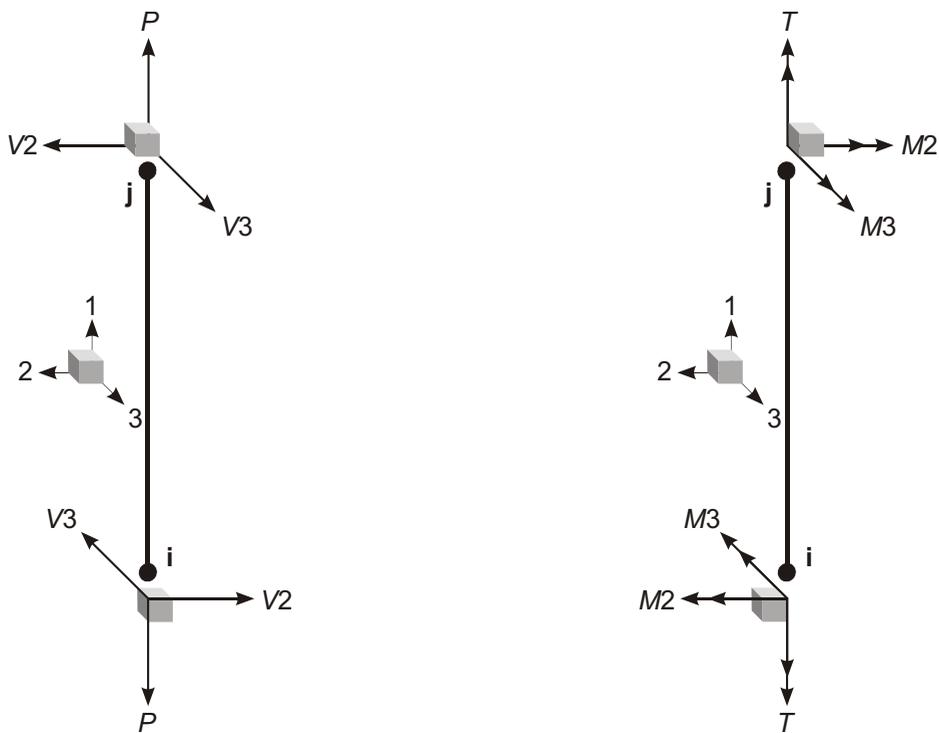


Figure 65

Link/Support Element Internal Forces and Moments, Shown Acting at the Joints

Spring Force-Deformation Relationships

There are six force-deformation relationships that govern the behavior of the element, one for each of the internal springs:

- Axial: f_{u1} vs. d_{u1}
- Shear: f_{u2} vs. d_{u2} , f_{u3} vs. d_{u3}
- Torsional: f_{r1} vs. d_{r1}
- Pure bending: f_{r2} vs. d_{r2} , f_{r3} vs. d_{r3}

where f_{u1} , f_{u2} , and f_{u3} are the internal-spring forces; and f_{r1} , f_{r2} , and f_{r3} are the internal-spring moments.

Each of these relationships may be zero, linear only, or linear/nonlinear for a given Link/Support Property. These relationships may be independent or coupled. The

forces and moments may be related to the deformation rates (velocities) as well as to the deformations.

Element Internal Forces

The Link/Support element internal forces, $P, V2, V3$, and the internal moments, $T, M2, M3$, have the same meaning as for the Frame element. These are illustrated in Figure 65 (page 266). These can be defined in terms of the spring forces and moments as:

- Axial: $P = f_{u1}$
- Shear in the 1-2 plane: $V2 = f_{u2}, \quad M3_s = (d - \mathbf{d}\mathbf{j}2) f_{u2}$
- Shear in the 1-3 plane: $V3 = f_{u3}, \quad M2_s = (d - \mathbf{d}\mathbf{j}3) f_{u3}$
- Torsion: $T = f_{r1}$
- Pure bending in the 1-3 plane: $M2_b = f_{r2}$
- Pure bending in the 1-2 plane: $M3_b = f_{r3}$

where d is the distance from joint \mathbf{j} . The total bending-moment resultants $M2$ and $M3$ composed of shear and pure-bending parts:

$$\begin{matrix} M2 & M2_s & M2_b \\ M3 & M3_s & M3_b \end{matrix}$$

These internal forces and moments are present at every cross section along the length of the element.

See Topic “Internal Force Output” (page 144) in Chapter “The Frame Element.”

Uncoupled Linear Force-Deformation Relationships

If each of the internal springs are linear and uncoupled, the spring force-deformation relationships can be expressed in matrix form as:

$$\begin{matrix} f_{u1} & k_{u1} & 0 & 0 & 0 & 0 & 0 & d_{u1} \\ f_{u2} & & k_{u2} & 0 & 0 & 0 & 0 & d_{u2} \\ f_{u3} & & & k_{u3} & 0 & 0 & 0 & d_{u3} \\ f_{r1} & & & & k_{r1} & 0 & 0 & d_{r1} \\ f_{r2} & & sym. & & & k_{r2} & 0 & d_{r2} \\ f_{r3} & & & & & & k_{r3} & d_{r3} \end{matrix} \quad (\text{Eqn. 1})$$

where k_{u1} , k_{u2} , k_{u3} , k_{r1} , k_{r2} , and k_{r3} are the linear stiffness coefficients of the internal springs.

This can be recast in terms of the element internal forces and displacements at joint **j** for a one-joint element as:

$$\begin{matrix}
 P \\
 V2 \\
 V3 \\
 T \\
 M2 \\
 M3 \text{ j}
 \end{matrix}
 \begin{matrix}
 k_{u1} & 0 & 0 & 0 & 0 & 0 \\
 & k_{u2} & 0 & 0 & 0 & \mathbf{dj2} k_{u2} \\
 & & k_{u3} & 0 & \mathbf{dj3} k_{u3} & 0 \\
 & & & k_{r1} & 0 & 0 \\
 \text{sym.} & & & & k_{r2} & \mathbf{dj3}^2 k_{u3} \\
 & & & & & k_{r3} & \mathbf{dj2}^2 k_{u2}
 \end{matrix}
 \begin{matrix}
 u_1 \\
 u_2 \\
 u_3 \\
 r_1 \\
 r_2 \\
 r_3 \text{ j}
 \end{matrix}
 \tag{Eqn. 2}$$

This relationship also holds for a two-joint element if all displacements at joint **I** are zero.

Similar relationships hold for linear damping behavior, except that the stiffness terms are replaced with damping coefficients, and the displacements are replaced with the corresponding velocities.

Consider an example where the equivalent shear and bending springs are to be computed for a prismatic beam with a section bending stiffness of EI in the 1-2 plane. The stiffness matrix at joint **j** for the 1-2 bending plane is:

$$\begin{matrix}
 V2 \\
 M3 \text{ j}
 \end{matrix}
 \begin{matrix}
 \frac{EI}{L^3} & 12 & 6L & u_2 \\
 & 6L & 4L^2 & r_3 \text{ j}
 \end{matrix}$$

From this it can be determined that the equivalent shear spring has a stiffness of $k_{u2} = 12 \frac{EI}{L^3}$ located at $\mathbf{dj2} = \frac{L}{2}$, and the equivalent pure-bending spring has a stiffness of $k_{r3} = \frac{EI}{L}$.

For an element that possesses a true moment hinge in the 1-2 bending plane, the pure-bending stiffness is zero, and $\mathbf{dj2}$ is the distance to the hinge. See Figure 66 (page 269).

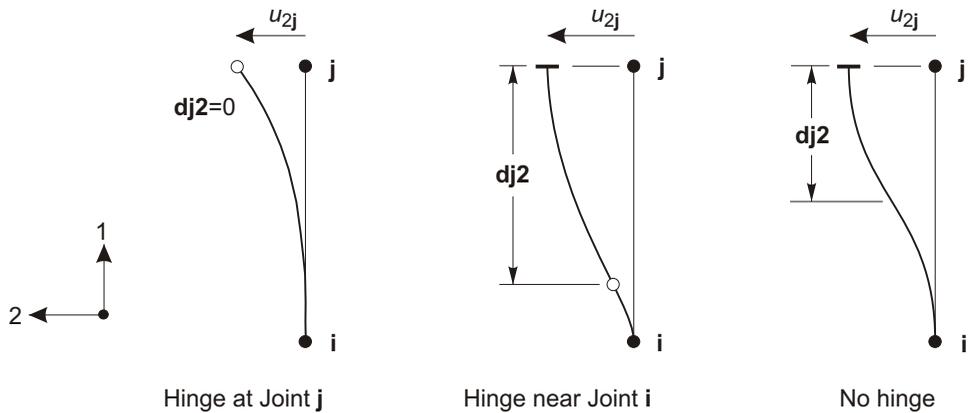


Figure 66

Location of Shear Spring at a Moment Hinge or Point of Inflection

Types of Linear/Nonlinear Properties

The primary Linear/Nonlinear Link/Support Properties may be of the following types:

- Coupled Linear
- Damper
- Gap
- Hook
- Multi-linear Elastic
- Multi-linear Plastic
- Plastic (Wen)
- Hysteretic (Rubber) Isolator
- Friction-Pendulum Isolator
- Tension/Compression Friction Pendulum Isolator

The first type, Coupled Linear, may have fully coupled linear stiffness and damping coefficients. This property type is described in Topic “Coupled Linear Property” (page 269) in this Chapter.

All other property types are considered nonlinear. However, for each nonlinear property type you also specify a set of uncoupled linear stiffness and damping coef-

ficients that are used instead of the nonlinear properties for linear analyses. These substitute linear properties are called “linear effective stiffness” and “linear effective damping” properties.

For more information:

- See Topic “Coupled Linear Property” (page 269) in this Chapter.
- See Chapter “The Link/Support Element—Advanced” (page 275).

Coupled Linear Property

The Coupled Linear Link/Support Property is fully linear. It has no nonlinear behavior. The linear behavior is used for all linear and nonlinear analyses. It is also used for frequency-dependent analyses unless frequency-dependent properties have been assigned to the Link/Support element.

The stiffness matrix of Eqn. (1) (page 267) may now be fully populated:

$$\begin{array}{cccccccc}
 f_{u1} & k_{u1} & k_{u1u2} & k_{u1u3} & k_{u1r1} & k_{u1r2} & k_{u1r3} & d_{u1} \\
 f_{u2} & & k_{u2} & k_{u2u3} & k_{u2r1} & k_{u2r2} & k_{u2r3} & d_{u2} \\
 f_{u3} & & & k_{u3} & k_{u3r1} & k_{u3r2} & k_{u3r3} & d_{u3} \\
 f_{r1} & & & & k_{r1} & k_{r1r2} & k_{r1r3} & d_{r1} \\
 f_{r2} & & sym. & & & k_{r2} & k_{r2r3} & d_{r2} \\
 f_{r3} & & & & & & k_{r3} & d_{r3}
 \end{array} \quad (\text{Eqn. 3})$$

where k_{u1} , k_{u1u2} , k_{u2} , k_{u1u3} , k_{u2u3} , k_{u3} , ..., k_{r3} are the linear stiffness coefficients of the internal springs.

The corresponding matrix of Eqn. (2) (page 268) can be developed from the relationships that give the element internal forces in terms of the spring forces and moments. See Topic “Element Internal Forces” (page 267) in this Chapter.

Similarly, the damping matrix is fully populated and has the same form as the stiffness matrix. Note that the damping behavior is active for all dynamic analyses. This is in contrast to linear effective damping, which is not active for nonlinear analyses.

Fixed Degrees of Freedom

It is possible to select any number or all of the element deformation degrees of freedom (DOF) to be fixed, i.e., that their deformation is zero. For a two-joint link, this

is equivalent to specifying a joint constraint. For a one-joint link, this is equivalent to specifying a restraint.

While this feature is convenient, it should be used with care. There are several important considerations that should be kept in mind.

- Link elements with fixed DOF should not be connected to other fixed link elements or connected to constrained joints. Otherwise, this will result in joints that are multiply-constrained, which may be inaccurate for dynamics.
- For link elements with non-zero length, fixing fewer than all six DOF can lead to loss of mass coupling between rotational and translational DOF, which can affect the dynamics of the structure if large masses or rotational inertias are present at the connected joints.
- When all six DOF are fixed, a full rigid-body constraint is used so that the translational and rotational inertia will be accurately handled, with no loss of coupling, provided that the link elements is not connected to other constrained joints or fixed links.
- Link forces are not reported for fixed DOF.

If forces are needed at fixed DOF, if fixed links are connected to constrained joints, or if inertial coupling will be lost for non-zero length elements, it is recommended that large stiffnesses be used rather than fixed DOF to represent rigid conditions. It is important to use stiffness values that are large enough to prevent significant deformation, but not so large as to cause numerical sensitivity in the equation solution, particularly near nonlinearities in the model. Choose stiffness values that are about 100 to 10000 times larger than those in connected elements that are expected to undergo deformation.

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. The mass contributed by the Link or Support element is lumped at the joints **I** and **j**. *No inertial effects are considered within the element itself.*

For each Link/Support Property, you may specify a total translational mass, **m**. Half of the mass is assigned to the three translational degrees of freedom at each of the element's one or two joints. For single-joint elements, half of the mass is assumed to be grounded.

You may additionally specify total rotational mass moments of inertia, **mr1**, **mr2**, and **mr3**, about the three local axes of each element. Half of each mass moment of

inertia is assigned to each of the element's one or two joints. For single-joint elements, half of each mass moment of inertia is assumed to be grounded.

The rotational inertias are defined in the element local coordinate system, but will be transformed by the program to the local coordinate systems for joint **I** and **j**. If the three inertias are not equal and element local axes are not parallel to the joint local axes, then cross-coupling inertia terms will be generated during this transformation. These will be discarded by the program, resulting in some error.

It is *strongly recommended* that there be mass corresponding to each nonlinear deformation load in order to generate appropriate Ritz vectors for nonlinear modal time-history analysis. Note that rotational inertia is needed as well as translational mass for nonlinear shear deformations if either the element length or **dj** is non-zero.

For more information:

- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Static and Dynamic Analysis” (page 341).
- See Topic “Nonlinear Deformation Loads” (page 271) in this Chapter.

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For each Link/Support Property, a total self-weight, **w**, may be defined. Half of this weight is assigned to each joint of each Link/Support element using that Link/Support Property. For single-joint elements, half of the weight is assumed to be grounded.

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns” for more information.

Gravity Load

Gravity Load can be applied to each Link/Support element to activate the self-weight of the element. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 272) in this Chapter for the definition of self-weight for the Link/Support element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Internal Force and Deformation Output

Link/Support element internal forces and deformations can be requested for Load Cases and Load Combinations.

Results for linear analyses are based upon the linear effective-stiffness and effective-damping properties and do not include any nonlinear effects. Only the results for nonlinear Load Cases include nonlinear behavior.

The element internal forces were defined in Subtopic “Element Internal Forces” (page 267) of this Chapter. The internal deformations were defined in Topic “Internal Deformations” (page 260) of this Chapter.

The element internal forces are labeled P, V2, V3, T, M2, and M3 in the output. The internal deformations are labeled U1, U2, U3, R1, R2, and R3 in the output, corresponding to the values of d_{u1} , d_{u2} , d_{u3} , d_{r1} , d_{r2} , and d_{r3} .

For more information:

- See Chapter “Load Patterns” (page 321).
- See Chapter “Load Cases” (page 341).

The Link/Support Element—Advanced

The basic, linear behavior of the Link and Support elements was described in the previous Chapter. The present Chapter describes the use of the Link and Support elements to model nonlinear behavior and frequency-dependent behavior.

Advanced Topics

- Overview
- Nonlinear Link/Support Properties
- Linear Effective Stiffness
- Linear Effective Damping
- Exponential Maxwell Damper Property
- Bilinear Maxwell Damper Property
- Friction-Spring Damper Property
- Gap Property
- Hook Property
- Wen Plasticity Property
- Multi-Linear Elastic Property
- Multi-Linear Plastic Property

- Hysteretic (Rubber) Isolator Property
- Friction-Pendulum Isolator Property
- Double-Acting Friction-Pendulum Isolator Property
- Triple-Pendulum Isolator Property
- Nonlinear Deformation Loads
- Frequency-Dependent Properties

Overview

The basic features of the Link and Support elements were described in the previous Chapter, “The Link/Support Element—Basic” (page 251).

This Chapter describes the various type of nonlinear properties that are available, the concepts of linear effective stiffness and damping, the use of nonlinear deformation loads for Ritz-vector analysis, and frequency-dependent properties.

Nonlinear Link/Support Properties

The **nonlinear properties** for each Link/Support Property must be of one of the various types described below. The type determines which degrees of freedom may be nonlinear and the kinds of nonlinear force-deformation relationships available for those degrees of freedom.

Every degree of freedom may have linear effective-stiffness and effective-damping properties specified, as described below in Subtopics “Linear Effective Stiffness” and “Linear Effective Damping.”

During *nonlinear analysis*, the nonlinear force-deformation relationships are used at all degrees of freedom for which nonlinear properties were specified. For all other degrees of freedom, the linear effective stiffnesses are used during a nonlinear analysis.

Linear analyses that start from zero initial conditions will use the linear effective stiffness regardless of whether nonlinear properties were specified or not. Linear analyses that use the stiffness from the end of a previous nonlinear analysis will use the nonlinear properties. Linear effective damping is used for all linear analyses, but it is not used for any nonlinear analysis.

Each nonlinear force-deformation relationship includes a stiffness coefficient, \mathbf{k} . This represents the linear stiffness when the nonlinear effect is negligible, e.g., for rapid loading of the Damper; for a closed Gap or Hook; or in the absence of yielding or slipping for the Plastic1, Isolator1, or Isolator2 properties. If \mathbf{k} is zero, no nonlinear force can be generated for that degree of freedom, with the exception of the pendulum force in the Isolator2 property.

IMPORTANT! You may sometimes be tempted to specify very large values for \mathbf{k} , particularly for Maxwell Dampers, Gap, and Hook properties. Resist this temptation! If you want to limit elastic deformations in a particular internal spring, it is usually sufficient to use a value of \mathbf{k} that is from 10^2 to 10^4 times as large as the corresponding stiffness in any connected elements. Larger values of \mathbf{k} may cause numerical difficulties during solution. See the additional discussion for the Maxwell Damper properties below.

Linear Effective Stiffness

For each nonlinear type of Link/Support Property, you may specify six uncoupled linear effective-stiffness coefficients, \mathbf{ke} , one for each of the internal springs.

The **linear effective stiffness** represents the total elastic stiffness for the Link/Support element that is used for *all linear analyses that start from zero initial conditions*. The actual nonlinear properties are ignored for these types of analysis.

If you do not specify nonlinear properties for a particular degree of freedom, then the linear effective stiffness is used for that degree of freedom for all linear and nonlinear analyses.

The effective force-deformation relationships for the Link/Support Properties are given by Equation 1 above with the appropriate values of \mathbf{ke} substituted for k_{u1} , k_{u2} , k_{u3} , k_{r1} , k_{r2} , and k_{r3} .

Special Considerations for Modal Analyses

The effective stiffness properties are not used for nonlinear degrees of freedom during nonlinear time-history analysis. However, nonlinear modal time-history (FNA) analyses do make use of the vibration modes that are computed based on the effective stiffness *if the modal analysis itself uses the stiffness from zero initial conditions*. During time integration the behavior of these modes is modified so that the structural response reflects the actual stiffness and other nonlinear parameters spec-

ified. The rate of convergence of the nonlinear iteration may be improved by changing the effective stiffness.

Following are some guidelines for selecting the linear effective stiffness. You should deviate from these as necessary to achieve your modeling and analysis goals. In particular, you should consider whether you are more interested in the results to be obtained from linear analyses, or in obtaining modes that are used as the basis for nonlinear modal time-history analyses.

- When carrying out analyses based on the UBC '94 code or similar, the effective stiffness should usually be the code-defined maximum effective stiffness
- For Gap and Hook elements the effective stiffness should usually be zero or \mathbf{k} , depending on whether the element is likely to be open or closed, respectively, in normal service
- For Damper elements, the effective stiffness should usually be zero
- For other elements, the stiffness should be between zero and \mathbf{k}
- If you have chosen an artificially large value for \mathbf{k} , be sure to use a much smaller value for \mathbf{ke} to help avoid numerical problems in nonlinear modal time-history analyses

In the above, \mathbf{k} is the nonlinear stiffness property for a given degree of freedom. See Chapter “The Link/Support Element—Basic” (page 251).

Remember that the above considerations *do not apply* for modes calculated using the stiffness from the end of a nonlinear static or nonlinear direct-integration time-history load case. In this case, the actual nonlinear stiffness of the links at the end of the nonlinear stiffness load case is used, and the effective stiffness is disregarded for nonlinear degrees of freedom.

For more information, see Topic “Nonlinear Modal Time-History Analysis (FNA)” (page 349) in Chapter “Nonlinear Time-History Analysis.”

Linear Effective Damping

For each nonlinear-type of Link/Support Property, you may specify six uncoupled linear effective-damping coefficients, \mathbf{ce} , one for each of the internal springs. By default, each coefficient \mathbf{ce} is equal to zero.

The **linear effective damping** represents the total viscous damping for the Link/Support element that is used for response-spectrum analyses, for linear and periodic time-history analyses, and for frequency-dependent analyses if fre-

quency-dependent properties have not been assigned to a given Link or Support element. The actual nonlinear properties are ignored for these types of analysis. Effective damping can be used to represent energy dissipation due to nonlinear damping, plasticity, or friction.

The effective force/deformation-rate relationships for the Link/Support Properties are given by Equation 1 above with the appropriate values of \mathbf{ce} substituted for k_{u1} , k_{u2} , k_{u3} , k_{r1} , k_{r2} , and k_{r3} , and deformation rates substituted for the corresponding deformations.

For response-spectrum and linear modal time-history analysis, the effective damping values are converted to modal damping ratios assuming proportional damping, i.e., the modal cross-coupling damping terms are ignored. These effective modal-damping values are added to any other modal damping that you specify directly. The program will not permit the total damping ratio for any mode to exceed 99.995%.

Important Note: Modal cross-coupling damping terms can be very significant for some structures. A linear analysis based on effective-damping properties may grossly overestimate or underestimate the amount of damping present in the structure.

Nonlinear time-history analysis is strongly recommended to determine the effect of added energy dissipation devices. Nonlinear time-history analysis does not use the effective damping values since it accounts for energy dissipation in the elements directly, and correctly accounts for the effects of modal cross-coupling.

Exponential Maxwell Damper Property

This element is very well suited for modeling viscous dampers that have a nonlinear force-velocity relationship. For simple linear damping, you may instead want to use the coupled linear Link Support Property. The linear property does not require the series spring used by the exponential Maxwell damper, and it does allow you to consider a parallel spring. See Topic “Coupled Linear Property” (page 270) for more information

For the exponential damper, you can specify independent damping properties for each deformational degree of freedom. The damping properties are based on the Maxwell model of viscoelasticity (Malvern, 1969) having a exponential viscous damper in series with a linear spring. See Figure 67 (page 280). If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

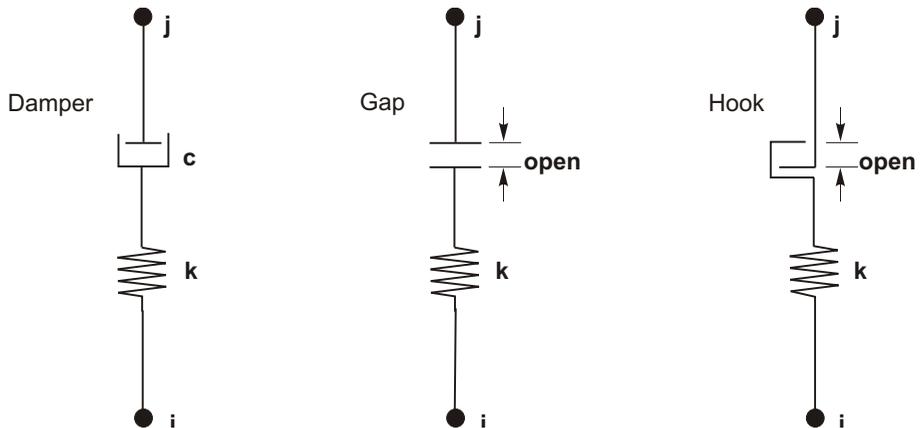


Figure 67
*Maxwell Dampers, Gap, and Hook Property Types,
 Shown for Axial Deformations*

The nonlinear force-deformation relationship is given by:

$$f = k d_k + c \dot{d}_c^{\text{cexp}}$$

where k is the spring constant, c is the damping coefficient, cexp is the damping exponent, d_k is the deformation across the spring, and \dot{d}_c is the deformation rate across the damper. The damping exponent must be positive; the practical range is between 0.2 and 2.0.

The spring and damping deformations sum to the total internal deformation:

$$d = d_k + d_c$$

The series spring is very important for capturing realistic behavior of nonlinear dampers, especially those with fractional exponents. It represents the elastic flexibility of the damping device, including the fluid column and the connecting mechanisms. It prevents the damping term from producing unrealistically large viscous forces at small velocities, which can have a *very significant* impact on overall structural behavior.

You may be tempted to introduce a large stiffness value, k , to represent “pure” damping, but this may result in unconservative and unrealistic behavior. It would be better to get a realistic value of the elastic flexibility from the manufacturer of

the device and the details of the connections, or make an engineering estimate of the value. For more information about the importance of the spring constant, see the *Software Verification Manual*, where SAP2000 results are compared with experiment.

Bilinear Maxwell Damper Property

This element is well suited for modeling viscous dampers, including oil dampers, that have a bilinear force-velocity relationship. Such dampers typically have a relief valve that reduces the damping coefficient when a certain relief (or “yield”) force level is reached. See Kasai (2004) for more information.

For the bilinear damper, you can specify independent damping properties for each deformational degree of freedom. The damping properties are based on the Maxwell model of viscoelasticity (Malvern, 1969) having a bilinear viscous damper in series with a linear spring. See Figure 67 (page 280). If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

The nonlinear force-deformation relationship is given by:

$$f = \mathbf{k} d_k + \begin{cases} \mathbf{c1} \dot{d}_c & \text{if } |\dot{d}_c| \leq \mathbf{f1} / \mathbf{c1} \\ \mathbf{f1} + \mathbf{c2} (\dot{d}_c - \mathbf{f1} / \mathbf{c1}) & \text{if } |\dot{d}_c| > \mathbf{f1} / \mathbf{c1} \end{cases}$$

where \mathbf{k} is the spring constant, $\mathbf{c1}$ is the damping coefficient for force levels below the relief force, $\mathbf{c2}$ is the damping coefficient for force levels above the relief force, $\mathbf{f1}$ is the relief force, d_k is the deformation across the spring, \dot{d}_c is the deformation rate across the damper, and $\mathbf{f1} / \mathbf{c1}$ is the yield deformation rate. The damping coefficients must satisfy the relationship $\mathbf{c1} > \mathbf{c2} \geq 0$.

The spring and damping deformations sum to the total internal deformation:

$$d = d_k + d_c$$

The series spring is very important for capturing realistic behavior of nonlinear dampers. It represents the elastic flexibility of the damping device, including the fluid column and the connecting mechanisms, which can have a *very significant* impact on overall structural behavior.

You may be tempted to introduce a large stiffness value, \mathbf{k} , to represent “pure” damping, but this may result in unconservative and unrealistic behavior. It would

be better to get a realistic value of the elastic flexibility from the manufacturer of the device and the details of the connections, or make an engineering estimate of the value. Expected values for oil dampers may be on the order of $\mathbf{c1} / \mathbf{k} = 0.1$ second.

The formulation used for this element works best if the time-step size is small compared to the period of excitation. This is usually not a problem for nonlinear modal (FNA) time-history analysis, which tends to automatically use very small time sub-steps. For nonlinear direct-integration time-history analysis, you should try different time-step sizes until consistent results are obtained. You can either change the size of the output time steps or change the maximum sub-step size (one of the nonlinear load-case parameters).

Friction-Spring Damper Property

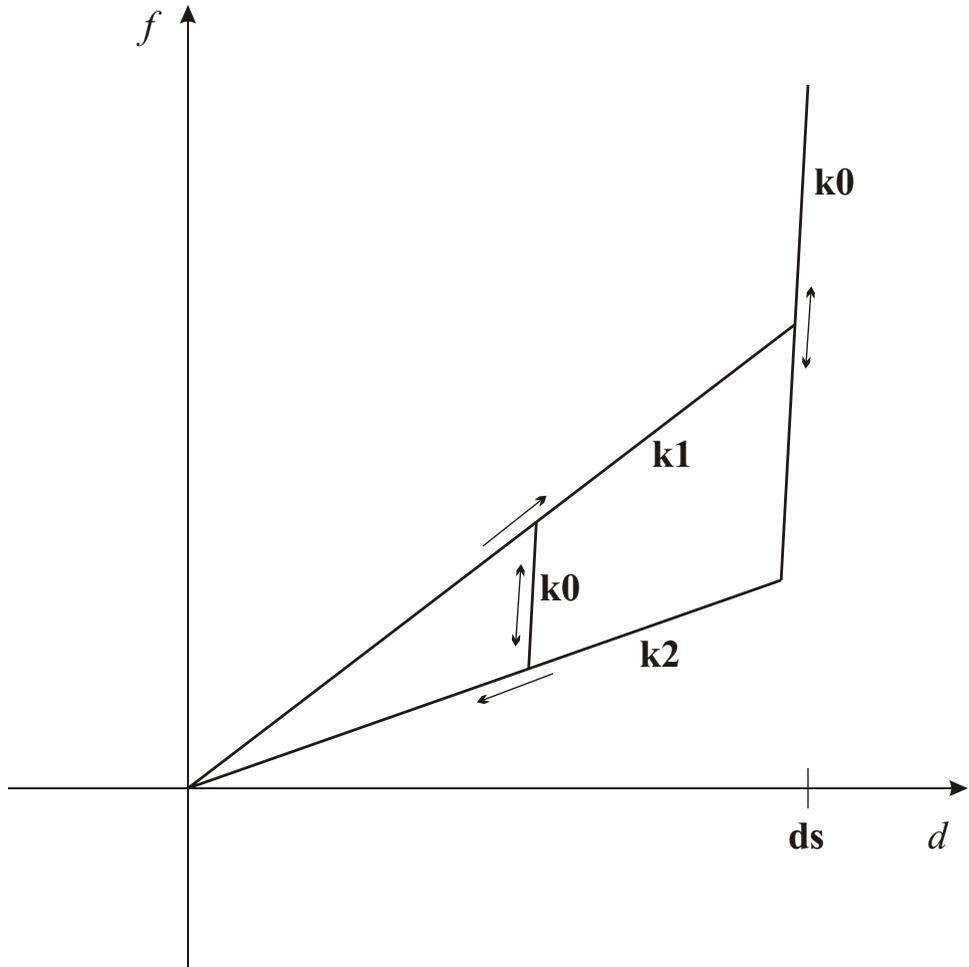
This element represents a hysteretic damping system consisting of a stack of alternating inner and outer metallic rings that exhibit internal friction when the stack is loaded in compression and again when the stack unloads. Each cycle of loading and unloading dissipates energy hysteretically, meaning that it is based on displacement and is independent of velocity.

Although the fundamental mechanism is compression of the ring stack, manufactured devices may exhibit external damping behavior in compression, tension, or both. In addition, kinematic mechanisms may be devised so that the frictional behavior can be applied to shear or moment degrees of freedom. For this reason, friction-spring behavior may be defined independently for any of the six degrees of freedom of a link element.

A friction-spring system under loading-slipping exhibits a linear stiffness $\mathbf{k1}$ that is a combination of the elastic stiffness to compress the inner rings and expand the outer rings, plus the frictional resistance. When unloading-slipping, the system exhibits a smaller linear stiffness $\mathbf{k2}$ that is a combination of the same elastic stiffness due to ring compression and expansion, but this time with the frictional resistance acting in the opposite direction. The elastic transition between loading and unloading stiffness is along a much larger elastic stiffness $\mathbf{k0}$.

This behavior is shown in Figure 68 (page 283). The stiffness values should satisfy the following relationship:

$$\mathbf{k0} \gg \mathbf{k1} > \mathbf{k2} > 0$$

**Figure 68**

Force-Deformation Behavior of a Friction-Spring Damper in Tension with Stop Displacement but No Precompression

There is usually a limiting stop displacement $\mathbf{ds} > 0$ beyond which the rings can no longer slip. All deformation beyond the stop displacement loads and unloads along the elastic stiffness of $\mathbf{k0}$. If you set $\mathbf{ds} = 0$ then no stop displacement is applied.

A precompression displacement $\mathbf{dc} < 0$ may also be specified. This has the effect of increasing the initial force at which slipping begins when loading to approximately $\mathbf{k1 dc}$, and at which slipping ends when unloading to approximately $\mathbf{k2 dc}$, as shown in Figure 69 (page 284). Loading and unloading from zero to these values

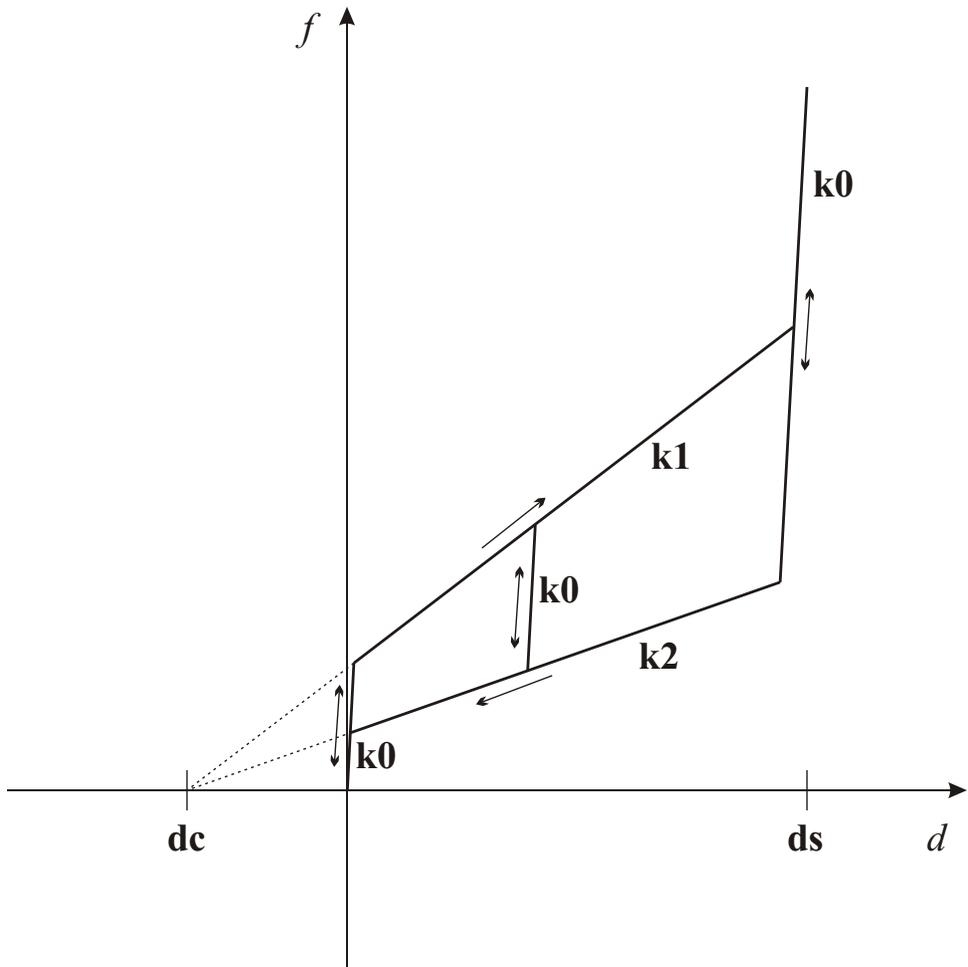
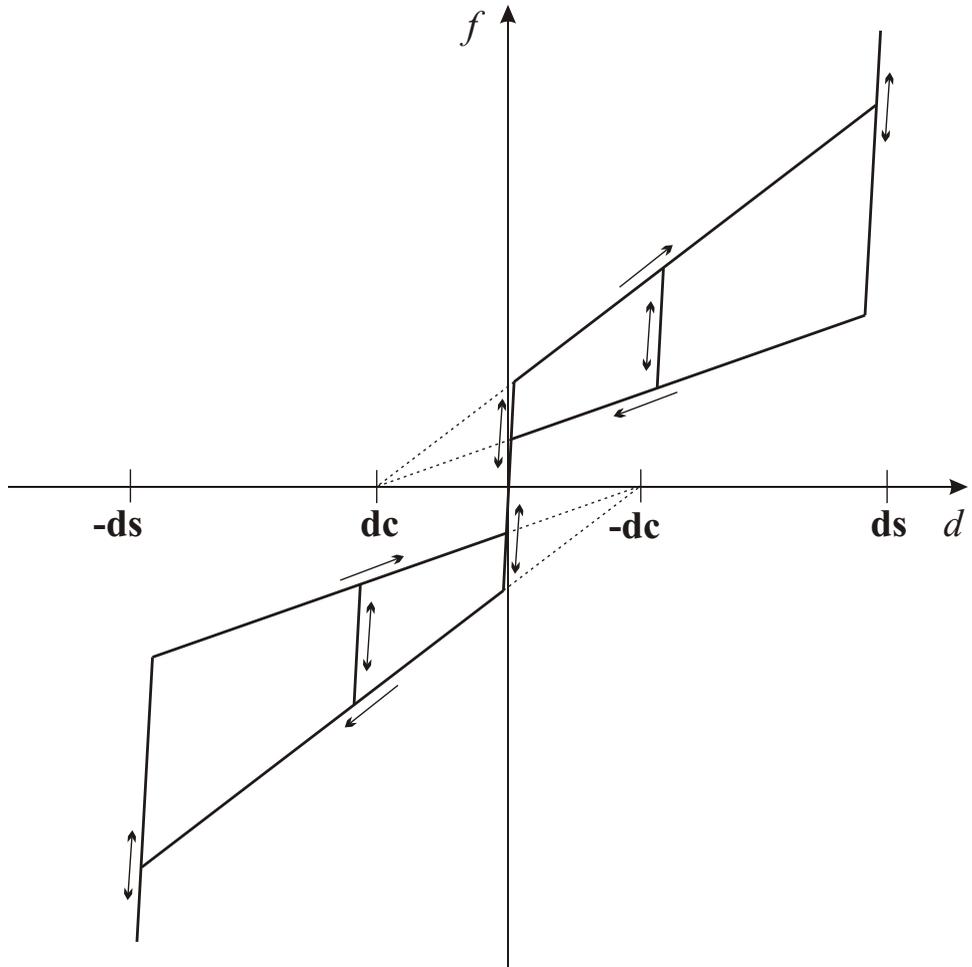


Figure 69

Force-Deformation Behavior of a Friction-Spring Damper in Tension with Stop Displacement and Precompression Displacement

occurs along the elastic stiffness of k_0 . If you set $dc = 0$ then no precompression displacement is applied, and slipping begins and ends at zero force.

The force-deformation relationships shown in Figure 68 and Figure 69 are for tension (or any positive force or moment). Similar behavior could be specified in compression as well. If tension and compression behavior are both requested, the behavior is symmetrical about the origin as shown in Figure 70 (page 285). If unsym-

**Figure 70**

Force-Deformation Behavior of a Friction-Spring Damper in Symmetrical Tension and Compression

metrical behavior is desired, you can use two elements in parallel, one for the tension behavior and one for the compression behavior.

Loading and unloading can occur along the solid lines, but only in the direction of the arrows. The dotted lines in Figure 69 and Figure 70 are shown for reference to the precompression displacement dc . However loading and unloading cannot occur along the dotted lines.

The nonlinear force-deformation relationship when slipping is given by:

$$f = k_1 (d - dc) \text{ when loading-slipping, } \dot{d} \dot{=} 0$$

$$f = k_2 (d - dc) \text{ when unloading-slipping, } \dot{d} \dot{=} 0$$

The force-displacement relationship is elastic along a line with stiffness **k0** in the transitions: (1) between sliding while loading and unloading, (2) from zero when **dc** < 0, and (3) after the stop displacement **ds** > 0.

The values of **k0**, **k1**, **k2**, **ds**, and **dc** should be obtained from the manufacturer of the device. The value of **k0** should be large compared to **k1** and **k2**, but not so large as to cause numerical sensitivity, which in turn may produce inaccurate results or nonlinear convergence problems. A realistic value is best, but it should probably be no more than about 10,000 times larger than **k1**.

Gap Property

For each deformational degree of freedom you may specify independent gap (“compression-only”) properties. See Figure 67 (page 280).

All internal deformations are independent. The opening or closing of a gap for one deformation does not affect the behavior of the other deformations.

If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

The nonlinear force-deformation relationship is given by:

$$f = \begin{cases} k (d - \text{open}) & \text{if } d - \text{open} \leq 0 \\ 0 & \text{otherwise} \end{cases}$$

where **k** is the spring constant, and **open** is the initial gap opening, which must be zero or positive.

Hook Property

For each deformational degree of freedom you may specify independent hook (“tension-only”) properties. See Figure 67 (page 280).

All internal deformations are independent. The opening or closing of a hook for one deformation does not affect the behavior of the other deformations.

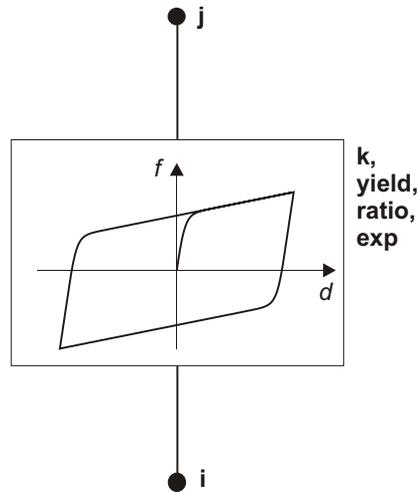


Figure 71
Wen Plasticity Property Type for Uniaxial Deformation

If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

The nonlinear force-deformation relationship is given by:

$$f = \begin{cases} k(d - \text{open}) & \text{if } d > \text{open} \\ 0 & \text{otherwise} \end{cases}$$

where **k** is the spring constant, and **open** is the initial hook opening, which must be zero or positive.

Wen Plasticity Property

For each deformational degree of freedom you may specify independent uniaxial-plasticity properties. The plasticity model is based on the hysteretic behavior proposed by Wen (1976). See Figure 71 (page 287).

All internal deformations are independent. The yielding at one degree of freedom does not affect the behavior of the other deformations.

If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

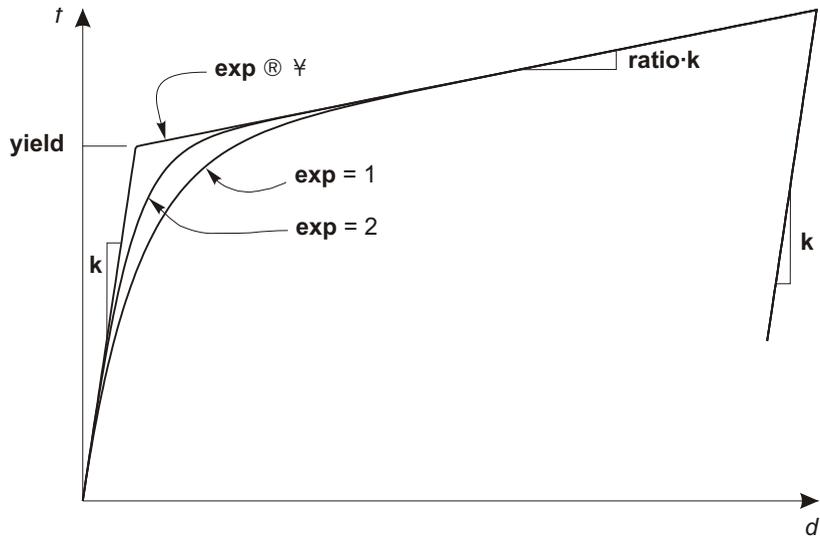


Figure 72
Definition of Parameters for the Wen Plasticity Property

The nonlinear force-deformation relationship is given by:

$$f = \text{ratio} \cdot k \cdot d \cdot (1 - \text{ratio}) \cdot \text{yield} \cdot z$$

where **k** is the elastic spring constant, **yield** is the yield force, **ratio** is the specified ratio of post-yield stiffness to elastic stiffness (**k**), and *z* is an internal hysteretic variable. This variable has a range of $|z| \leq 1$, with the yield surface represented by $|z| = 1$. The initial value of *z* is zero, and it evolves according to the differential equation:

$$\dot{z} = \frac{k}{\text{yield}} \cdot \frac{\dot{d} (1 - |z|^{\text{exp}})}{\dot{d}} \quad \text{if } \dot{d} z > 0$$

$$\dot{z} = 0 \quad \text{otherwise}$$

where **exp** is an exponent greater than or equal to unity. Larger values of this exponent increases the sharpness of yielding as shown in Figure 72 (page 288). The practical limit for **exp** is about 20. The equation for \dot{z} is equivalent to Wen's model with $A = 1$ and $\beta = 0.5$.

Multi-Linear Elastic Property

For each deformational degree of freedom you may specify multi-linear elastic properties.

All internal deformations are independent. The deformation in one degree of freedom does not affect the behavior of any other. If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

The nonlinear force-deformation relationship is given by a multi-linear curve that you define by a set of points. The curve can take on almost any shape, with the following restrictions:

- One point must be the origin, (0,0)
- At least one point with positive deformation, and one point with negative deformation, must be defined
- The deformations of the specified points must increase monotonically, with no two values being equal
- The forces (moments) can take on any value

The slope given by the last two specified points on the positive deformation axis is extrapolated to infinite positive deformation. Similarly, the slope given by the last two specified points on the negative deformation axis is extrapolated to infinite negative deformation.

The behavior is nonlinear but it is elastic. This means that the element loads and unloads along the same curve, and no energy is dissipated. On the other hand, the Multi-Linear Plastic Property defined next does exhibit hysteresis and dissipates energy under reverse and cyclic loading.

See also Topic “Hysteresis Models” (page 85) in Chapter “Material Properties.”

Multi-Linear Plastic Property

This model is similar to the Multi-Linear Elastic Property defined above, but the Plastic property exhibits different loading and unloading behavior, and thus dissipates energy according to various hysteresis models that are available. For each deformational degree of freedom you may specify independent multi-linear plasticity properties.

All internal deformations are independent. The deformation in one degree of freedom does not affect the behavior of any other. If you do not specify nonlinear properties for a degree of freedom, that degree of freedom is linear using the effective stiffness, which may be zero.

The nonlinear force-deformation relationship is given by a multi-linear curve that you define by a set of points. The curve can take on almost any shape, with the following restrictions:

- One point must be the origin, (0,0)
- At least one point with positive deformation, and one point with negative deformation, must be defined
- The deformations of the specified points must increase monotonically, with no two values being equal
- The forces (moments) at a point must have the same sign as the deformation (they can be zero)
- The slope given by the last two points specified on the positive deformation axis is extrapolated to infinite positive deformation, or until it reaches zero value. Similarly, the slope given by the last two points specified on the negative deformation axis is extrapolated to infinite negative deformation, or until it reaches zero value.

The given curve defines the force-deformation (moment-rotation) relationship under monotonic loading. The first slope on either side of the origin is elastic; the remaining segments define plastic deformation. If the deformation reverses after plastic deformation, unloading occurs along a different, steeper curve before beginning plastic deformation in the reverse direction.

This process of loading and unloading along different curves dissipates energy and is called hysteresis. Several different hysteresis models are available to describe the behavior of different types of materials. For the most part, these differ in the amount of energy they dissipate in a given cycle of deformation, and how the energy dissipation behavior changes with an increasing amount of deformation.

The available hysteresis models for the Multi-Linear Plastic Property are described in Topic “Hysteresis Models” (page 85) in Chapter “Material Properties”. These include kinematic, degrading, Takeda, pivot, concrete, BRB hardening, and isotropic.

Hysteretic (Rubber) Isolator Property

This is a biaxial hysteretic isolator that has coupled plasticity properties for the two shear deformations, and linear effective-stiffness properties for the remaining four deformations. The plasticity model is based on the hysteretic behavior proposed by Wen (1976), and Park, Wen and Ang (1986), and recommended for base-isolation analysis by Nagarajaiah, Reinhorn and Constantinou (1991). See Figure 73 (page 292).

For each *shear* deformation degree of freedom you may independently specify either linear or nonlinear behavior:

- If both shear degrees of freedom are nonlinear, the coupled force-deformation relationship is given by:

$$\begin{aligned} f_{u2} &= \text{ratio2 } k2 d_{u2} + (1 - \text{ratio2}) \text{yield2 } z_2 \\ f_{u3} &= \text{ratio3 } k3 d_{u3} + (1 - \text{ratio3}) \text{yield3 } z_3 \end{aligned}$$

where **k2** and **k3** are the elastic spring constants, **yield2** and **yield3** are the yield forces, **ratio2** and **ratio3** are the ratios of post-yield stiffnesses to elastic stiffnesses (**k2** and **k3**), and z_2 and z_3 are internal hysteretic variables. These variables have a range of $\sqrt{z_2^2 + z_3^2} \leq 1$, with the yield surface represented by $\sqrt{z_2^2 + z_3^2} = 1$. The initial values of z_2 and z_3 are zero, and they evolve according to the differential equations:

$$\begin{aligned} \dot{z}_2 &= 1 - a_2 z_2^2 - a_3 z_2 z_3 \frac{k2}{\text{yield2}} \dot{d}_{u2} \\ \dot{z}_3 &= a_2 z_2 z_3 + 1 - a_3 z_3^2 \frac{k3}{\text{yield3}} \dot{d}_{u3} \end{aligned}$$

Where:

$$\begin{aligned} a_2 &= \begin{cases} 1 & \text{if } \dot{d}_{u2} z_2 < 0 \\ 0 & \text{otherwise} \end{cases} \\ a_3 &= \begin{cases} 1 & \text{if } \dot{d}_{u3} z_3 < 0 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

These equations are equivalent to those of Park, Wen and Ang (1986) with $A = 1$ and $\alpha = 0.5$.

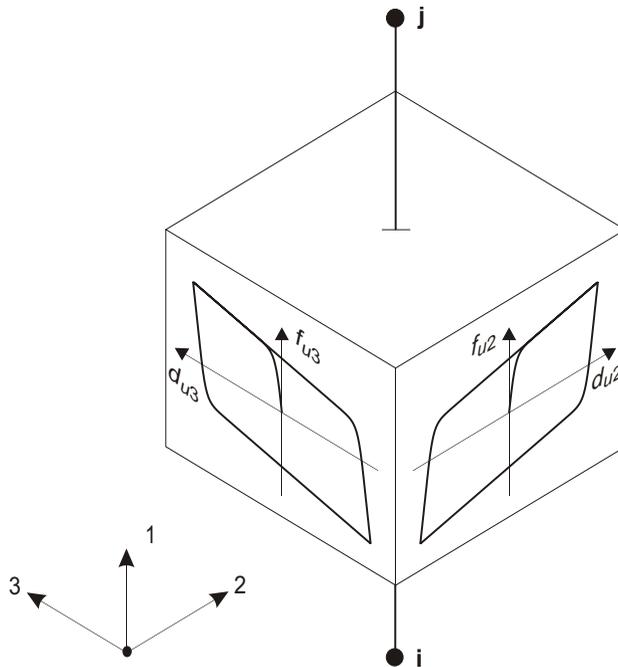


Figure 73
Hysteretic Isolator Property for Biaxial Shear Deformation

- If only one shear degree of freedom is nonlinear, the above equations reduce to the uniaxial plasticity behavior of the Plastic1 property with **exp 2** for that degree of freedom.

A linear spring relationship applies to the axial deformation, the three moment deformations, and to any shear deformation without nonlinear properties. All linear degrees of freedom use the corresponding effective stiffness, which may be zero.

Friction-Pendulum Isolator Property

This is a biaxial friction-pendulum isolator that has coupled friction properties for the two shear deformations, post-slip stiffness in the shear directions due the pendulum radii of the slipping surfaces, gap behavior in the axial direction, and linear effective-stiffness properties for the three moment deformations. See Figure 74 (page 293).

$$f_{u1} = P \begin{cases} \mathbf{k1} d_{u1} & \text{if } d_{u1} < 0 \\ 0 & \text{otherwise} \end{cases}$$

In order to generate nonlinear shear force in the element, the stiffness $\mathbf{k1}$ must be positive, and hence force P must be negative (compressive). A reasonable value for the stiffness $\mathbf{k1}$ can be obtained as the AE/L of the device, and should include the flexibility of the connections or supports that is not otherwise included in the model.

You may additionally specify a damping coefficient, $\mathbf{c1}$, for the axial degree of freedom, in which case the axial force becomes:

$$f_{u1} = P \begin{cases} \mathbf{c1} \dot{d}_{u1} & \text{if } d_{u1} < 0 \\ 0 & \text{otherwise} \end{cases}$$

The damping force only exists when the isolator is in compression, regardless of the sign of the velocity.

Force f_{u1} is the total axial force exerted by the element on the connected joints. However, only the stiffness force P is assumed to act on the bearing surface, causing shear resistance. The damping force is external.

The purpose of the damping coefficient is to reduce the numerical chatter (oscillation) that can be present in some analyses. You can estimate the damping coefficient needed to achieve a certain ratio, r , of critical damping (e.g., $r = 0.05$) from the formula

$$r = \frac{\mathbf{c1}}{2\sqrt{\mathbf{k1} m}}$$

where m is the tributary mass for the isolator, which could be estimated from the self-weight axial force divided by the acceleration due to gravity. It is up to you to verify the applicability of this approach for your particular application. See the *Software Verification Manual* for a discussion on the use of this damping coefficient.

Shear Behavior

For each *shear* deformation degree of freedom you may independently specify either linear or nonlinear behavior:

- If both shear degrees of freedom are nonlinear, the friction and pendulum effects for each shear deformation act in parallel:

$$\begin{matrix} f_{u2} & f_{u2f} & f_{u2p} \\ f_{u3} & f_{u3f} & f_{u3p} \end{matrix}$$

The frictional force-deformation relationships are given by:

$$\begin{matrix} f_{u2f} & P_{u2} z_2 \\ f_{u3f} & P_{u3} z_3 \end{matrix}$$

where μ_2 and μ_3 are friction coefficients, and z_2 and z_3 are internal hysteretic variables. The friction coefficients are velocity-dependent according to:

$$\begin{matrix} \mu_2 & \mathbf{fast2} & (\mathbf{fast2} & \mathbf{slow2}) e^{-r v} \\ \mu_3 & \mathbf{fast3} & (\mathbf{fast3} & \mathbf{slow3}) e^{-r v} \end{matrix} \quad (\text{Eqn. 1a})$$

where **slow2** and **slow3** are the friction coefficients at zero velocity, **fast2** and **fast3** are the friction coefficients at fast velocities, v is the resultant velocity of sliding:

$$v = \sqrt{\dot{d}_{u2}^2 + \dot{d}_{u3}^2} \quad (\text{Eqn. 1b})$$

r is an effective inverse velocity given by:

$$r = \frac{\mathbf{rate2} \dot{d}_{u2}^2 + \mathbf{rate3} \dot{d}_{u3}^2}{v^2} \quad (\text{Eqn. 1c})$$

and **rate2** and **rate3** are the inverses of characteristic sliding velocities. For a Teflon-steel interface the coefficient of friction normally increases with sliding velocity (Nagarajaiah, Reinhorn, and Constantinou, 1991).

The internal hysteretic variables have a range of $\sqrt{z_2^2 + z_3^2} \leq 1$, with the yield surface represented by $\sqrt{z_2^2 + z_3^2} = 1$. The initial values of z_2 and z_3 are zero, and they evolve according to the differential equations:

$$\begin{matrix} \dot{z}_2 & = & 1 - a_2 z_2^2 - a_3 z_2 z_3 & \frac{\mathbf{k2}}{P_{u2}} \dot{d}_{u2} \\ \dot{z}_3 & = & a_2 z_2 z_3 - 1 + a_3 z_3^2 & \frac{\mathbf{k3}}{P_{u3}} \dot{d}_{u3} \end{matrix}$$

where k_2 and k_3 are the elastic shear stiffnesses of the slider in the absence of sliding, and

$$a_2 = \begin{cases} 1 & \text{if } \dot{d}_{u2} z_2 > 0 \\ 0 & \text{otherwise} \end{cases}$$

$$a_3 = \begin{cases} 1 & \text{if } \dot{d}_{u3} z_3 > 0 \\ 0 & \text{otherwise} \end{cases}$$

These equations are equivalent to those of Park, Wen and Ang (1986) with $A = 1$ and $\alpha = 0.5$.

This friction model permits some sliding at all non-zero levels of shear force; the amount of sliding becomes much larger as the shear force approaches the “yield” value of P . Sliding at lower values of shear force can be minimized by using larger values of the elastic shear stiffnesses. However, realistic values of the shear stiffness are recommended, and can be estimated as AG/L of the locked-up device. This should also include the flexibility of the connections or supports that is not otherwise included in the model.

The pendulum force-deformation relationships are given by:

$$f_{u2p} = P \frac{d_{u2}}{\text{radius2}}$$

$$f_{u3p} = P \frac{d_{u3}}{\text{radius3}}$$

The specified radius should actually be the effective pendulum length, which is the radius of the sliding surface minus the distance from the surface to the point of articulation of the device. This information should be obtained from the manufacturer of the device.

Normally the radii in the two shear directions will be equal (spherical surface), or one radius will be zero (cylindrical surface). However, it is permitted to specify unequal non-zero radii.

A zero radius indicates a flat surface, and the corresponding perpendicular shear force is zero. This can be used to model general gap-friction behavior.

- If only one shear degree of freedom is nonlinear, the above frictional equations reduce to:

$$f_f = P z$$

$$\dot{z} = \frac{\mathbf{k}}{P} \dot{d} (1 - z^2) \quad \text{if } \dot{d} z \geq 0$$

$$0 \quad \text{otherwise}$$

The above pendulum equation is unchanged for the nonlinear degree of freedom.

Linear Behavior

A linear spring relationship applies to the three moment deformations, and to any shear deformation without nonlinear properties. All linear degrees of freedom use the corresponding effective stiffness, which may be zero. The axial degree of freedom is always nonlinear for nonlinear analyses.

Double-Acting Friction-Pendulum Isolator Property

This is a biaxial friction-pendulum isolator that supports tension as well as compression, and has uncoupled behavior in the two shear directions. The frictional resistance can be different depending on whether the isolator is in tension or compression. This device consists of two orthogonal, curved rails that are interlocked together. It is intended to provide seismic isolation with uplift prevention, and is described in detail by Roussis and Constantinou [2005].

Axial Behavior

Independent stiffnesses and gap openings may be specified for tension and compression. The axial force, f_{u1} , is *always* nonlinear, and is given by:

$$f_{u1} = P \begin{cases} \mathbf{k1c}(d_{u1} - \mathbf{openc}) & \text{if } (d_{u1} - \mathbf{openc}) \leq 0 \\ \mathbf{k1t}(d_{u1} - \mathbf{opent}) & \text{if } (d_{u1} - \mathbf{opent}) > 0 \\ 0 & \text{otherwise} \end{cases}$$

where **k1c** is the compressive stiffness, **k1t** is the tensile stiffness, **openc** is the gap opening in compression, and **opent** is the gap opening in tension. Each of the four values may be zero or positive.

You may additionally specify a damping coefficient, **c1**, for the axial degree of freedom, in which case the axial force becomes:

$$f_{u1} = P + c1 \dot{d}_{u1}$$

The damping force exists whether the isolator is in tension, compression, or is gapping.

Force f_{u1} is the total axial force exerted by the element on the connected joints. However, only the stiffness force P is assumed to act on the bearing surface, causing shear resistance. The damping force is external. See Topic “Friction-Pendulum Isolator Property” (page 292) for a discussion on the use of this damping.

Shear Behavior

For each shear deformation degree of freedom you may independently specify either linear or nonlinear behavior. The behavior in the two shear directions is uncoupled, although they both depend on the same axial force P .

For each nonlinear shear degree of freedom $u2$ or $u3$, you independently specify the following parameters:

- Stiffness **k**, representing the elastic behavior before sliding begins. This value is the same for positive or negative P .
- Friction coefficients **slowc** and **fastc** for friction under compression at different velocities, and coefficients **slowt** and **fastt** for friction under tension at different velocities.
- Rate parameters **ratec** and **ratet** for friction under compression and tension, respectively. These are the inverses of characteristic sliding velocities. For a Teflon-steel interface the coefficient of friction normally increases with sliding velocity (Nagarajaiah, Reinhorn, and Constantinou, 1991).
- Radius **radius**, which is the same for tension and compression.

Looking at one shear direction, and considering either tension or compression using the appropriate friction parameters, the shear force f is given by:

$$f = f_f + f_p$$

$$f_f = P \cdot z$$

$$f_p = (\text{fast} - (\text{fast} - \text{slow}) e^{-\text{rate} \dot{d}})$$

$$\dot{z} = \frac{k}{P} \dot{d} (1 - z^2) \quad \text{if } \dot{d} z > 0$$

$$\dot{z} = 0 \quad \text{otherwise}$$

$$f_p = P \frac{d}{\text{radius}}$$

where d is the shear deformation and z is an internal hysteretic variable. In the above, the indicators for shear degree of freedom u_2 or u_3 , as well as for tension or compression, have been dropped.

Linear Behavior

A linear spring relationship applies to the three moment deformations, and to any shear deformation without nonlinear properties. All linear degrees of freedom use the corresponding effective stiffness, which may be zero. The axial degree of freedom is always nonlinear for nonlinear analyses.

Triple-Pendulum Isolator Property

This is a biaxial friction-pendulum isolator (or bearing) that has four spherical sliding surfaces, two outer and two inner, as shown in Figure 75 (page 301). Kinematic considerations couple the motion of the two inner sliders, resulting in behavior that can be modeled as three friction-pendulum isolators in series. By choosing different friction coefficients and radii for the various sliding surfaces, the transverse (shear) force-deflection curve can be tailored to provide different response for different levels of excitation. The properties of the two inner sliders are assumed to be identical. The characteristics and performance of this type of isolator are described in Morgan and Mahin (2008).

Each of the three pendula that act series for the triple pendulum uses the same mathematical model as described in the topic “Friction-Pendulum Isolator Property” (page 292), with differences as described below. The axial and moment behavior are for the isolator as a whole. The shear forces are equal in the three pendula, while the shear deformations and velocities are additive.

The friction forces and pendulum forces are directly proportional to the compressive axial force in the element. The element cannot carry axial tension.

Axial Behavior

The axial force, f_{u1} , is *always* nonlinear, and is given by:

$$f_{u1} = P \begin{cases} \mathbf{k1} d_{u1} & \text{if } d_{u1} \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

In order to generate nonlinear shear force in the element, the stiffness $\mathbf{k1}$ must be positive, and hence force P must be negative (compressive).

You may additionally specify a damping coefficient, $\mathbf{c1}$, for the axial degree of freedom, in which case the axial force becomes:

$$f_{u1} = \begin{cases} P + \mathbf{c1} \dot{d}_{u1} & \text{if } d_{u1} < 0 \\ 0 & \text{otherwise} \end{cases}$$

The damping force only exists when the isolator is in compression.

Force f_{u1} is the total axial force exerted by the element on the connected joints. However, only the stiffness force P is assumed to act on the sliding surfaces, causing shear resistance. The damping force is external.

The purpose of the damping coefficient is to reduce the numerical chatter (oscillation) that can be present in some analyses. You can estimate the damping coefficient needed to achieve a certain ratio, r , of critical damping (e.g., $r = 0.05$) from the formula

$$r = \frac{\mathbf{c1}}{2\sqrt{\mathbf{k1} m}}$$

where m is the tributary mass for the isolator, which could be estimated from the self-weight axial force divided by the acceleration due to gravity. It is up to you to verify the applicability of this approach for your particular application. See the *Software Verification Manual* for a discussion on the use of this damping coefficient for the friction-pendulum isolator.

Shear Behavior

For each *shear* deformation degree of freedom you may independently specify either linear or nonlinear behavior. In the usual case where both the $u2$ and $u3$ degrees of freedom are nonlinear, the behavior is isotropic and the properties are identical in both directions. In other words, the sliding surfaces are spherical.

The four sliding surfaces are identified as:

- Inner Bottom (index 1)
- Inner Top, identical to Inner Bottom (index 1)
- Outer Bottom (index 2)
- Outer Top (index 3)

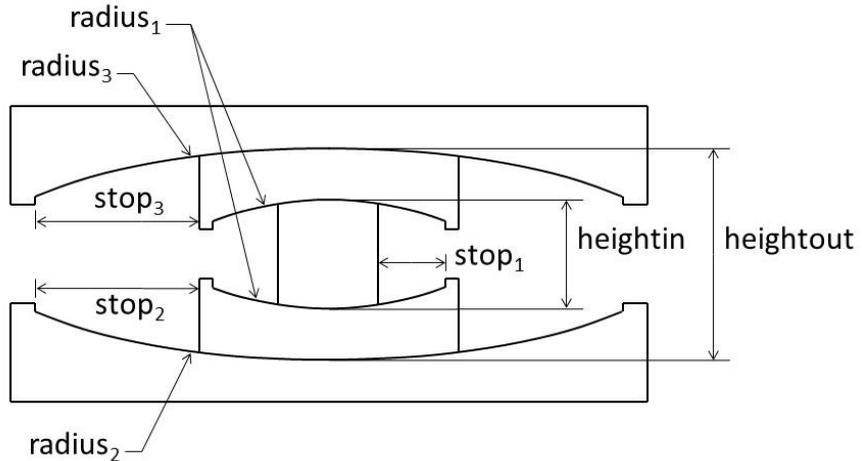


Figure 75
Triple-Pendulum Isolator with Four Sliding Surfaces
The two inner surfaces are always symmetrical

These are shown in Figure 75 (page 301). In the discussion below, the indices shown above will be used to identify the three pendulum mechanisms.

For each sliding surface, the following isotropic properties are specified:

- **stiff**, shear stiffness before sliding, and stiffness at the stop
- **slow**, friction coefficient at zero velocity
- **fast**, friction coefficient at fast velocity
- **rate**, inverse of the characteristic velocity for friction
- **radius**, actual spherical radius of the surface
- **stop**, transverse displacement permitted before encountering a stiff stop

In addition, two heights are specified giving the distances between the sliding surfaces:

- **heightin**, maximum distance between the two inner surfaces at zero displacement
- **heightout**, maximum distance between the two outer surfaces at zero displacement

The isolator stiffness at each stage of sliding is dependent upon the effective pendulum lengths. For each surface, the pendulum length is given by the radius minus the distance to the point of articulation, a fraction of the distance between the surfaces. Thus:

$$L_1 = \text{radius}_1 - 0.5 \text{ heightin}$$

$$L_2 = \text{radius}_2 - \text{heightout}$$

$$L_3 = \text{radius}_3 - (1 - \text{heightout})$$

where

$$\text{heightout} = \frac{\text{radius}_2}{\text{radius}_2 + \text{radius}_3}$$

During sliding, the net friction coefficient for each surface is dependent upon the velocity of sliding on that surface according to the rate-dependent friction equations (Eqns. 1, page 293) given for the friction-pendulum isolator.

The monotonic, radial force-deflection behavior of a typical triple-pendulum isolator is shown in Figure 76 (page 303). The isolator is assumed to be in compression, and the plotted shear force, \tilde{V} , is normalized by the compression load: $\tilde{V} = |V/P|$, $P > 0$.

The actual behavior will depend on the relative pendulum lengths, friction coefficients, and stop distances. For the purpose of the following discussion, we will assume that the net coefficients for the three surfaces are constant and ordered such that $\mu_1 < \mu_2 < \mu_3$, and that the stop distances are large enough as described. The following behavior is observed:

- $0 < u < u_a$: The initial resistance is given by the elastic stiffness until first sliding begins. This occurs at a normalized shear force equal to the lowest friction coefficient, μ_1 , for the inner sliding surfaces.
- $u_a < u < u_b$: The shear stiffness is now governed by the pendulum lengths of the two inner surfaces, both of which are sliding. This stiffness is $1/(2L_1)$.
- $u_b < u < u_c$: When the shear force increases to the second friction coefficient, μ_2 , sliding shifts from the inner-bottom surface to the outer-bottom surface, and the stiffness becomes $1/(L_1 + L_2)$.
- $u_c < u < u_d$: Similarly, at the level of the third friction coefficient, μ_3 , sliding shifts from the inner-top surface to the outer-top surface, and the stiffness becomes $1/(L_2 + L_3)$.

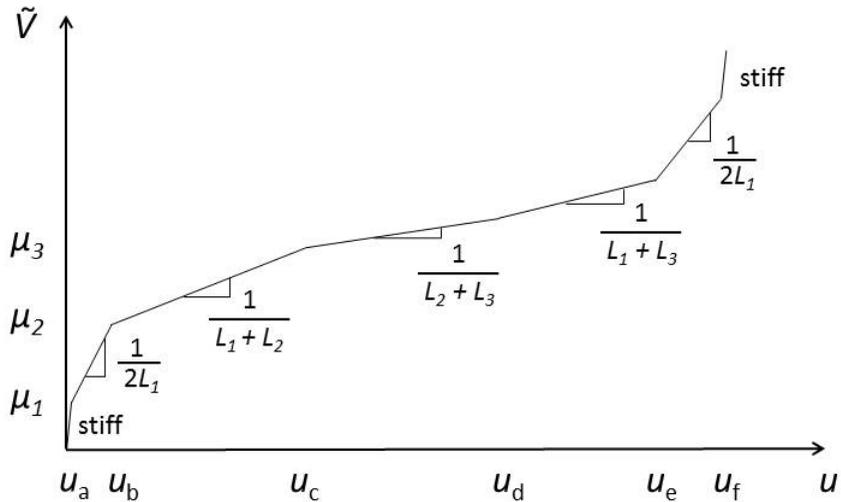


Figure 76

*Normalized Force-Deflection Curve for the Triple-Pendulum Isolator
Typical behavior for monotonic radial loading at slow speed*

- $u_d - u_e$: With increased sliding, one of the stops is hit, say for the outer-bottom surface. Sliding shifts back to the adjacent inner surface, and the stiffness becomes $1/(L_1 - L_3)$.
- $u_e - u_f$: When the outer-top stop is hit, sliding shifts back to the adjacent inner surface and the stiffness becomes $1/(2L_1)$ again.
- $u - u_f$: After both inner stops are hit, the initial elastic stiffness is engaged.

The actual behavior will depend on the relative values of the parameters used for the actual isolator, the velocities of sliding as they affect the friction coefficients, and the history and direction of the motion.

The transition to sliding is not abrupt, but rather occurs gradually as the normalized shear force approaches each friction coefficient. This friction model permits some sliding at all non-zero levels of shear force; the amount of sliding becomes much larger as the shear force approaches the “yield” value of P . Sliding at lower values of shear force can be minimized by using larger values of the elastic shear stiffnesses, **stiff**. However, realistic values of the shear stiffness are recommended, and can be estimated as AG/L of the locked-up device. This should also include the flexibility of the connections or supports that is not otherwise included in the

model. In no case should the stiffness values be more than about 100 to 10000 times larger than $|P/L|$, where P is a typical axial compression value for the isolator, and L is the pendulum length. Larger values serve no practical purpose and make nonlinear convergence more difficult.

Linear Behavior

A linear spring relationship applies to the three moment deformations, and to any shear deformation without nonlinear properties. All linear degrees of freedom use the corresponding effective stiffness, which may be zero. The axial degree of freedom is always nonlinear for nonlinear analyses.

Nonlinear Deformation Loads

A **nonlinear deformation load** is a set of forces and/or moments on the structure that activates a nonlinear internal deformation of an Link/Support element. A nonlinear internal deformation is an Link/Support internal deformation for which nonlinear properties have been specified.

Nonlinear deformation loads are used as starting load vectors for Ritz-vector analysis. Their purpose is to generate Modes that can adequately represent nonlinear behavior when performing nonlinear *modal* time-history analyses. A separate nonlinear deformation load should be used for each nonlinear internal deformation of each Link/Support element.

When requesting a Ritz-vector analysis, you may specify that the program use built-in nonlinear deformation loads, or you may define your own Load Patterns for this purpose. In the latter case you may need up to six of these Load Patterns per Link/Support element in the model.

The built-in nonlinear deformation loads for a single two-joint Link element are shown in Figure 77 (page 305). Each set of forces and/or moments is self-equilibrating. This tends to localize the effect of the load, usually resulting in a better set of Ritz-vectors. For a single-joint element, only the forces and/or moments acting on joint j are needed.

It is *strongly recommended* that mass or mass moment of inertia be present at each degree of freedom that is acted upon by a force or moment from a nonlinear deformation load. This is needed to generate the appropriate Ritz vectors.

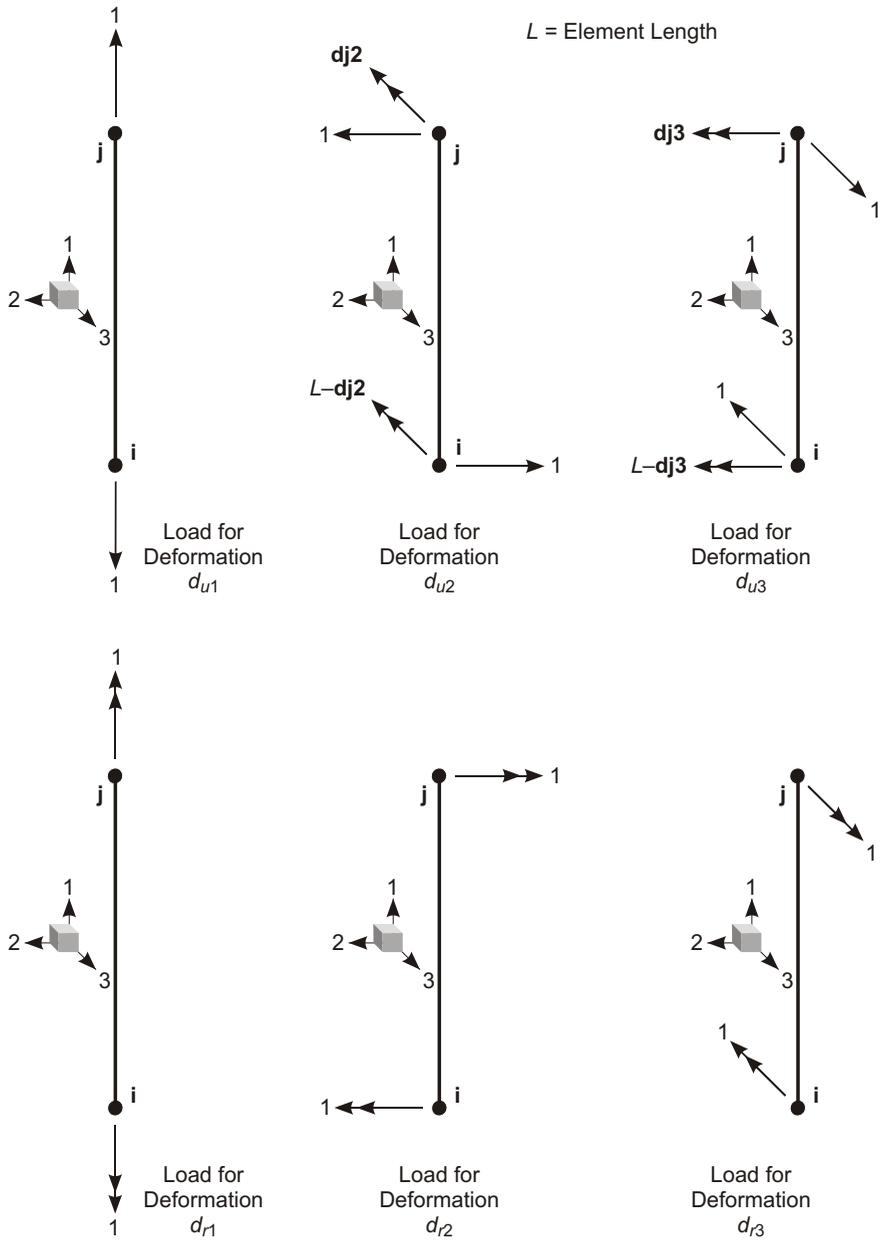


Figure 77
Built-in Nonlinear Deformation Loads for a Two-joint Link Element

For more information:

- See Topic “Internal Deformations” (page 260) in this Chapter.
- See Topic “Link/Support Properties” (page 263) in this Chapter.
- See Topic “Mass” (page 271) in this Chapter.
- See Topic “Ritz-Vector Analysis” (page 349) in Chapter “Load Cases.”
- See Topic “Nonlinear Modal Time-History Analysis (FNA)” (page 349) in Chapter “Load Cases.”

Frequency-Dependent Link/Support Properties

Each Link or Support element can have an optional set of frequency-dependent properties assigned to it in addition to the linear/nonlinear property set that must always be assigned. Frequency-dependent properties are only used for frequency-domain types of analyses, such as Steady-State or Power-Spectral-Density analyses.

Frequency-dependent properties represent the complex **impedance** of the element. There is a real part that represents the stiffness and inertial effects, and an imaginary part that represents the hysteretic damping effects. Frequency-dependent properties for the six degrees of freedom of the element may be coupled or uncoupled, as given by:

$$\begin{array}{cccccccc}
 f_{u1} & z_{u1} & z_{u1u2} & z_{u1u3} & z_{u1r1} & z_{u1r2} & z_{u1r3} & d_{u1} \\
 f_{u2} & & z_{u2} & z_{u2u3} & z_{u2r1} & z_{u2r2} & z_{u2r3} & d_{u2} \\
 f_{u3} & & & z_{u3} & z_{u3r1} & z_{u3r2} & z_{u3r3} & d_{u3} \\
 f_{r1} & & & & z_{r1} & z_{r1r2} & z_{r1r3} & d_{r1} \\
 f_{r2} & & \text{sym.} & & & z_{r2} & z_{r2r3} & d_{r2} \\
 f_{r3} & & & & & & z_{r3} & d_{r3}
 \end{array} \quad (\text{Eqn. 2})$$

where z_{u1} k_{u1} ic_{u1} is the impedance term in the $u1$ degree of freedom, and where k_{u1} is the stiffness/inertial component, c_{u1} is the damping component, and i is the square root of -1 . The other impedance terms are similar.

In Eqn. (2), the force terms on the left-hand side of the equation and the displacement terms on the right-hand side of the equations are also complex. The real parts of these terms represent the behavior at a phase angle of zero, with time variation given by the cosine function, and the imaginary parts represent behavior at a phase angle of 90° , with time variation given by the sine function.

Each of the 21 impedance terms may vary with frequency. You define the variation for each term as a set of points giving stiffness vs. frequency and damping vs. frequency. It is not unusual for the stiffness term to be negative over part of the range.

A common use for frequency-dependent properties would be in Support elements that represent the far-field radiation-damping effect of the soil region under a rigid foundation.

The Tendon Object

Tendons are a special type of object that can be embedded inside other objects (frames, shells, planes, asolids, and solids) to represent the effect of prestressing and post-tensioning. These tendons attach to the other objects and impose load upon them.

Advanced Topics

- Overview
- Geometry
- Discretization
- Tendons Modeled as Loads or Elements
- Connectivity
- Degrees of Freedom
- Local Coordinate Systems
- Section Properties
- Nonlinear Properties
- Mass
- Prestress Load

- Self-Weight Load
- Gravity Load
- Temperature Load
- Strain Load
- Deformation Load
- Target-Force Load
- Internal Force Output

Overview

Tendons are a special type of object that can be embedded inside other objects (frames, shells, planes, asolids, and solids) to represent the effect of prestressing and post-tensioning. These tendons attach to the other objects through which they pass and impose load upon them.

You may specify whether the tendons are to be modeled as independent elements in the analysis, or just to act upon the rest of the structure as loads. Modeling as loads is adequate for linear analyses when you know the losses that will be caused by elastic shortening and time-dependent effects.

Tendons should be modeled as elements if you want the program to calculate the losses due to elastic shortening and time-dependent effects, if you want to consider nonlinearity in the Tendons, or if you want to know the forces acting in the Tendons due to other loading on the structure.

Tendon objects share some features with Frame elements, which will be cross-referenced in this Chapter.

Geometry

Any number of tendons may be defined. Each tendon is drawn or defined as a type of line object between two joints, **I** and **j**. The two joints must not share the same location in space. The two ends of the Tendon are denoted end I and end J, respectively.

The Tendon may have an arbitrary curved or segmented shape in three dimensions between those points, and may be offset at the ends from these joints.

Discretization

A Tendon may be a long object with complicated geometry, but it will be automatically discretized into shorter segments for purposes of analysis. You must specify the maximum length of these discretization segments during the definition of the Tendon. These lengths can affect how the Tendon loads the structure and the accuracy of the analysis results. You should choose shorter lengths for Tendons with highly curved geometry, or Tendons that pass through parts of the structure with complicated geometry or changes in properties. If you are not sure what value to use, try several different values to see how they affect the results.

Tendons Modeled as Loads or Elements

You have a choice for each Tendon how it is to be modeled for analysis:

- As equivalent loads that act upon the structure
- As independent elements with stiffness, mass and loading

Modeling as loads is adequate for linear analyses when you know in advance the losses that will be caused by elastic shortening and time-dependent effects.

Tendons should be modeled as elements if you want the program to calculate the losses due to elastic shortening and time-dependent effects, if you want to consider nonlinearity in the Tendons, or if you want to know the forces acting in the Tendons due to other loading on the structure. The discretized Tendon is internally analyzed as a series of equivalent short, straight Frame elements.

Connectivity

The Tendon is connected to Frame, Shell, Plane, Asolid, and Solid elements through which it passes along its length. This connection is made automatically by the program. See additional description below for parts of the tendon which do not fall inside another element.

To determine the elements through which the Tendon passes, the program uses the concept of a **bounding box**:

- For Frame elements, the bounding box is a rectangular prism bounded by the length of the element and its maximum cross-sectional dimensions in the local 2 and 3 directions.

- For Shell, Plane, and Asolid elements, it is the hexahedron bounded by the four sides of the element and the upper and lower surfaces in the local 3 direction, with thickness being considered.
- For Solid elements, it is the volume bounded by the six faces.

By default, the Tendon will be checked for connection against all elements in the model. You may restrict this by specifying a group of objects to which the Tendon may connect. The Tendon will not connect to any objects that are not in that group. See Topic “Groups” (page 9) in Chapter “Objects and Elements” for more information. Only elements from objects in this group are considered in the discussion below, where they are called “**bounding elements.**”

For Tendons modeled as loads, if any portion of the Tendon passes through a bounding element, load from that portion of the tendon within the bounding element is transferred to the joints of that element in a statically equivalent fashion. The load on any portion of a Tendon that does not fall within a bounding element is transferred to the nearest joint that is connected to a bounding element.

For Tendons modeled as elements, if any discretization point (i.e., either end of a discretization segment) falls within a bounding element, that point is connected by an interpolation constraint to all joints of that element. This means that for large discretizations, the tendon may not actually be connected to every element through which it passes. If the end of the entire Tendon object does not fall within a bounding element, it is connected by constraint to the nearest joint that is connected to a bounding element. However, internal discretization points that do not fall within a bounding element will not be connected to any element (except to adjacent tendon elements themselves), so that the tendon will be external to the structure at those locations.

Degrees of Freedom

The Tendon object has six degrees of freedom along its length. However, its effect upon the structure depends upon the elements to which it connects. When connecting to Frame and Shell elements, it may transmit forces and moments to the joints in those elements. When connecting to Planes, Asolids, and Solids, it only transmits forces to the joints.

Even when modeled as elements, a Tendon adds no additional degrees of freedom to a structure, since it is always constrained to act with the elements that contain it. The exception would be if there is a portion of the Tendon which is not embedded in any other element and acts as a external length of tendon.

For more information, please see Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom.”

Local Coordinate Systems

Each Tendon object has two local coordinate systems:

- **Base-line local coordinate system**, which is fixed for the whole object
- **Natural local coordinate system**, which varies along the length of the Tendon

These are described in the following.

Base-line Local Coordinate System

The Tendon base-line local coordinate system is used only to define the Tendon natural local coordinate system.

The axes of base-line system are denoted 1, 2 and 3. The first axis is directed along the straight line connecting the joints **I** and **j** that were used to define the Tendon. The remaining two axes lie in the plane perpendicular to this axis with an orientation that you specify. The base-line local coordinate system is fixed for the length of the Tendon, regardless of the Tendon’s trajectory in space.

Base-line local axes are defined exactly the same as for a Frame element connected to joints **I** and **j**, except the Tendon has zero joint offsets. Please see Topics “Local Coordinate System” (page 108) and “Advanced Local Coordinate System” (page 110) in Chapter “The Frame Element”.

Natural Local Coordinate System

The Tendon natural local coordinate system is used to define section properties, loads, and internal force output. This coordinate system is defined with respect to the base-line local coordinate system as follows:

- The 1 direction is directed along the tangent to the Tendon, in the direction from end **I** to end **J**.
- The 2 direction is parallel to the 1-2 plane of the base-line local coordinate system.
- The 3 direction is computed as the cross product of the natural local 1 and 2 directions.

See Topic “Local Coordinate Systems” (page 312) in this Chapter for more information.

Section Properties

A **Tendon Section** is a set of material and geometric properties that describe the cross-section of one or more Tendon objects. Sections are defined independently of the Tendons, and are assigned to the Tendon objects.

The cross section shape is always circular. The Section has axial, shear, bending and torsional properties, although we are primarily interested in only the axial behavior.

Material Properties

The material properties for the Section are specified by reference to a previously-defined Material. Uniaxial or isotropic material properties are used, even if the Material selected was defined as orthotropic or anisotropic. The material properties used by the Section are:

- The modulus of elasticity, **e1**, for axial stiffness and bending stiffness
- The shear modulus, **g12**, for torsional stiffness and transverse shear stiffness
- The coefficient of thermal expansion, **a1**, for axial expansion and thermal bending strain
- The mass density, **m**, for computing element mass
- The weight density, **w**, for computing Self-Weight Loads

The material properties **e1**, **g12**, and **a1** are all obtained at the material temperature of each individual Tendon object, and hence may not be unique for a given Section.

See Chapter “Material Properties” (page 69) for more information.

Geometric Properties and Section Stiffnesses

The cross section shape is always circular. You may specify either the diameter or the area, **a**. The axial stiffness of the Section is given by **a e1**.

The remaining section properties are automatically calculated for the circular shape. These, along with their corresponding Section stiffnesses, are given by:

- The moment of inertia, i_{33} , about the 3 axis for bending in the 1-2 plane, and the moment of inertia, i_{22} , about the 2 axis for bending in the 1-3 plane. The corresponding bending stiffnesses of the Section are given by $i_{33} e_1$ and $i_{22} e_1$;
- The torsional constant, j . The torsional stiffness of the Section is given by $j g_{12}$. For a circular section, the torsional constant is the same as the polar moment of inertia.
- The shear areas, as_2 and as_3 , for transverse shear in the 1-2 and 1-3 planes, respectively. The corresponding transverse shear stiffnesses of the Section are given by $as_2 g_{12}$ and $as_3 g_{12}$.

Note that although six stiffness values are computed, only the axial stiffness is of significance in a practical model. The other terms provide stability but have little influence on the behavior of the model.

Tension/Compression Limits

You may specify a maximum tension and/or a maximum compression that a Tendon may take. In the most common case, you would define **no-compression** behavior by specifying the compression limit to be zero. Even this may not be needed for bonded tendons, which can take some compression.

If you specify a tension limit, it must be zero or a positive value. If you specify a compression limit, it must be zero or a negative value. If you specify a tension and compression limit of zero, the Tendon will carry no axial force.

The tension/compression limit behavior is elastic. Any axial extension beyond the tension limit or axial shortening beyond the compression limit will occur with zero axial stiffness. These deformations are recovered elastically at zero stiffness.

Important! Tension/compression limits are nonlinear, and only affect Tendons that are modeled as elements, not Tendons modeled as loads. Furthermore, they are only effective in a nonlinear analysis.

Linear analyses starting from zero conditions (the unstressed state) behave as if there were no tension/compression limits. Linear analyses that use the stiffness from the end of a previous nonlinear analysis will have a constant stiffness: zero if a tension or compression limit was exceeded at the end of the nonlinear case, otherwise the elastic material stiffness of the tendon.

Mass

In a dynamic analysis, the mass of the structure is used to compute inertial forces. When modeled as elements, the mass contributed by the Tendon is lumped at each discretization point along the length of the Tendon. When modeled as loads, no mass is contributed to the model. This is not usually of any significance since the mass of a Tendon is generally small.

The total mass of the Tendon is equal to the integral along the length of the mass density, \mathbf{m} , multiplied by the cross-sectional area, \mathbf{a} .

For more information, see Topic “Mass Source” (page 334) in Chapter “Load Patterns”.

Prestress Load

Each Tendon produces a set of self-equilibrating forces and moments that act on the rest of the structure. You may assign different Prestress loading in different Load Patterns.

In a given Load Pattern, the Prestress Load for any Tendon is defined by the following parameters:

- Tension in the Tendon, before losses.
- Jacking location, either end I or end J, where the tensioning of the Tendon will occur
- Curvature coefficient. This specifies the fraction of tension loss (due to friction) per unit of angle change (in radians) along the length of the Tendon, measured from the jacking end.
- Wobble coefficient. This specifies the fraction of tension loss (due to friction) per unit of Tendon length, measured from the jacking end, due to imperfect straightness of the tendon.
- Anchorage set slip. This specifies the length of slippage at the jacking end of the Tendon due to the release of the jacking mechanism.

The following additional load parameters may be specified that only apply when the Tendon is modeled as loads:

- Elastic shortening stress, due to compressive shortening in the elements that are loaded by the Tendon. This may be due to loads from the Tendon itself or from other loads acting on the structure.

- Creep stress, due to compressive creep strains in the elements that are loaded by the Tendon.
- Shrinkage stress, due to compressive shrinkage strains in the elements that are loaded by the Tendon.
- Steel relaxation stress, due to tensile relaxation strains in the Tendon itself.

For Tendons modeled as elements, the elastic shortening stress is automatically accounted for in all analyses; the time-dependent creep, shrinkage, and relaxation stresses can be accounted for by performing a time-dependent staged-construction analysis. See Topic “Staged Construction” (page 439) in Chapter “Nonlinear Static Analysis” for more information.

The tension and all of the loss effects are converted to an equivalent set of distributed and concentrated forces acting along the length of the Tendon. These can be displayed and tabulated in the graphical user interface. They are applied to the model as follows:

- For Tendons modeled as loads, these forces are transferred to the joints of the bounding elements as described earlier in this chapter under Topic “Connectivity”
- For Tendons modeled as elements, the tension at each end of a discretized tendon element is converted to an equivalent strain load, $P/(a \ e1)$, where P is the local tension after losses, and may vary along the length.

To account for complicated jacking procedures, you can specify different prestress loads in different Load Patterns and apply them as appropriate.

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. For a Tendon object, the self-weight is a force that is distributed along the length of the element. The magnitude of the self-weight is equal to the weight density, w , multiplied by the cross-sectional area, a .

Self-Weight Load always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies equally to all elements in the structure.

For more information:

- See Topic “Weight Density” (page 78) in Chapter “Material Properties” for the definition of w .

- See Topic “Section Properties” (page 314) in this Chapter for the definition of **a**.
- See Topic “Self-Weight Load” (page 325) in Chapter “Load Patterns.”

Gravity Load

Gravity Load can be applied to each Tendon to activate the self-weight of the object. Using Gravity Load, the self-weight can be scaled and applied in any direction. Different scale factors and directions can be applied to each element.

If all elements are to be loaded equally and in the downward direction, it is more convenient to use Self-Weight Load.

For more information:

- See Topic “Self-Weight Load” (page 134) in this Chapter for the definition of self-weight for the Frame element.
- See Topic “Gravity Load” (page 326) in Chapter “Load Patterns.”

Temperature Load

Temperature Load creates thermal strain in the Tendon object. This strain is given by the product of the Material coefficient of thermal expansion and the temperature change of the object. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis.

For any Load Pattern, you may specify a Load Temperature field that is constant over the cross section and produces axial strains. This temperature field may be constant along the element length or interpolated from values given at the joints.

See Chapter “Load Patterns” (page 321) for more information.

Strain Load

Axial strain load represents a change in length per unit length. Positive strain increases the length of an unrestrained element, or causes compression in a restrained element. Strain and temperature load act similarly, differing only by a scale factor, namely the coefficient of thermal expansion.

Strain loads may be constant along the element length or linearly interpolated from values given at the joints by a Joint Pattern.

For more information, see Topic “Internal Force Output” (page 144) in this chapter, and also Chapter “Load Patterns” (page 321.)

Deformation Load

While Strain Load specifies a change in deformation per unit length, Deformation Load specifies the total deformation between the two ends of an unrestrained element. Deformation Load is internally converted to Strain Load, so you should choose whichever type of loading is most conveniently specified for your particular application. The specified axial deformation is converted to axial Strain Load by simply dividing by the element length. The computed strain loads are assumed to be constant along the length of the element.

See Chapter “Load Patterns” (page 321) for more information.

Target-Force Load

Target-Force Load is a special type of loading where you specify a desired axial force, and deformation load is iteratively applied to achieve the target force. Since the axial force may vary along the length of the element, you must also specify the relative location where the desired force is to occur. Target-Force loading is only used for nonlinear static and staged-construction analysis. If applied in any other type of Load Case, it has no effect.

Unlike all other types of loading, target-force loading is not incremental. Rather, you are specifying the total force that you want to be present in the frame element at the end of the Load Case or construction stage. The applied deformation that is calculated to achieve that force may be positive, negative, or zero, depending on the force present in the element at the beginning of the analysis. When a scale factor is applied to a Load Pattern that contains Target-Force loads, the total target force is scaled. The increment of applied deformation that is required may change by a different scale factor.

See Topic “Target-Force Load” (page 331) in Chapter “Load Patterns” and Topic “Target-Force Iteration” (page 444) in Chapter “Nonlinear Static Analysis” for more information.

Internal Force Output

The **Tendon internal forces** are the axial forces, P , that result from integrating the axial stresses over the object cross section. These internal forces are present at every cross section along the length of the Tendon, and may be plotted or tabulated as part of the analysis results. As is always true in SAP2000, tension is positive.

Important! Internal force output is only available for Tendons that are modeled as elements.

Load Patterns

A **Load Pattern** is a specified spatial distribution of forces, displacements, temperatures, and other effects that act upon the structure. A Load Pattern by itself does not cause any response of the structure. Load Patterns must be applied in Load Cases in order to produce results.

Basic Topics for All Users

- Overview
- Load Patterns, Load Cases, and Load Combinations
- Defining Load Patterns
- Coordinate Systems and Load Components
- Force Load
- Restraint Displacement Load
- Spring Displacement Load
- Self-Weight Load
- Concentrated Span Load
- Distributed Span Load
- Tendon Prestress Load

- Uniform Load
- Mass Source
- Acceleration Loads

Advanced Topics

- Gravity Load
- Surface Pressure Load
- Pore Pressure Load
- Temperature Load
- Strain and Deformation Load
- Rotate Load
- Joint Patterns

Overview

Each Load Pattern may consist of an arbitrary combination of the available load types:

- Concentrated forces and moments acting at the joints
- Displacements at the grounded ends of joint restraints, joint springs, and one-joint Link/Support elements
- Self-weight and/or gravity acting on all element types
- Concentrated or distributed forces and moments acting on the Frame elements
- Distributed forces acting on the Shell elements
- Surface pressure acting on the Shell, Plane, Asolid, and Solid elements
- Pore pressure acting on the Plane, Asolid, and Solid elements
- Thermal expansion acting on the Frame, Shell, Plane, Asolid, and Solid elements
- Prestress load due to Tendons acting in Frame, Shell, Plane, Asolid, and Solid elements
- Centrifugal forces acting on Asolid elements

For practical purposes, it is usually most convenient to restrict each Load Pattern to a single type of load, using Load Cases and Load Combinations to create more complicated combinations.

Load Patterns, Load Cases, and Load Combinations

A **Load Pattern** is a specified spatial distribution of forces, displacements, temperatures, and other effects that act upon the structure. A Load Pattern by itself does not cause any response of the structure.

Load Patterns must be applied in **Load Cases** in order to produce results. A Load Case defines how the Load Patterns are to be applied (e.g., statically or dynamically), how the structure responds (e.g., linearly or nonlinearly), and how the analysis is to be performed (e.g., modally or by direct-integration.) A Load Case may apply a single Load Pattern or a combination of Load Patterns.

The results of Load Cases can be combined after analysis by defining **Load Combinations**, also called **Combos**. A Load Combination is a sum or envelope of the results from different Load Cases. For linear problems, algebraic-sum types of Load Combinations make sense. For nonlinear problems, it is usually best to combine Load Patterns in the Load Cases, and use Load Combinations only for computing envelopes.

When printing, plotting, or displaying the response of the structure to loads, you may request results for Load Cases and Load Combinations, but not directly for Load Patterns.

When performing design, only the results from Load Combinations are used. Load Combinations can be automatically created by the design algorithms, or you can create your own. If necessary, you can define Load Combinations that contain only a single Load Case.

- See Chapter “Load Cases” (page 341).
- See Topic “Load Combinations (Combos)” (page 351) in Chapter “Load Cases”.

Defining Load Patterns

You can define as many Load Patterns as you want, each with a unique name that you specify. Within each Load Pattern, any number of joints or elements may be loaded by any number of different load types.

Each Load Pattern has a design type, such as DEAD, WIND, or QUAKE. This identifies the type of load applied so that the design algorithms know how to treat the load when it is applied in a Load Case.

Coordinate Systems and Load Components

Certain types of loads, such as temperature and pressure, are scalars that are independent of any coordinate system. Forces and displacements, however, are vectors whose components depend upon the coordinate system in which they are specified.

Vector loads may be specified with respect to any fixed coordinate system. The fixed coordinate system to be used is specified as **csys**. If **csys** is zero (the default), the global system is used. Otherwise **csys** refers to an Alternate Coordinate System.

The X, Y, and Z components of a force or translation in a fixed coordinate system are specified as **ux**, **uy**, and **uz**, respectively. The X, Y, and Z components of a moment or rotation are specified as **rx**, **ry**, and **rz**, respectively.

Most vector loads may also be specified with respect to joint and element local coordinate systems. Unlike fixed coordinate systems, the local coordinate systems may vary from joint to joint and element to element.

The 1, 2, and 3 components of a force or translation in a local coordinate system are specified as **u1**, **u2**, and **u3**, respectively. The 1, 2, and 3 components of a moment or rotation are specified as **r1**, **r2**, and **r3**, respectively.

You may use a different coordinate system, as convenient, for each application of a given type of load to a particular joint or element. The program will convert all these loads to a single coordinate system and add them together to get the total load.

See Chapter “Coordinate Systems” (page 11) for more information.

Effect upon Large-Displacements Analysis

In a large-displacements analysis, all loads specified in a joint or element local coordinate system will rotate with that joint or element. All loads specified in a fixed coordinate system will not change direction during the analysis.

For linear analyses, and analyses considering only P-delta geometric nonlinearity, the direction of loading does not change during the analysis.

See Chapter “Geometric Nonlinearity” (page 409) for more information.

Force Load

Force Load applies concentrated forces and moments to the joints. You may specify components **ux**, **uy**, **uz**, **rx**, **ry**, and **rz** in any fixed coordinate system **csys**, and components **u1**, **u2**, **u3**, **r1**, **r2**, and **r3** in the joint local coordinate system. Force values are additive after being converted to the joint local coordinate system.

See Topic “Force Load” (page 42) in Chapter “Joints and Degrees of Freedom” for more information.

Ground Displacement Load

Ground Displacement Load applies specified ground displacements (translations and rotations) along the supported degrees of freedom at the joints. Supports include joint restraints, joint springs, and one-joint Link/Support elements. You may specify components **ux**, **uy**, **uz**, **rx**, **ry**, and **rz** in any fixed coordinate system **csys**, and components **u1**, **u2**, **u3**, **r1**, **r2**, and **r3** in the joint local coordinate system. Displacement values are additive after being converted to the joint local coordinate system.

See Topic “Ground Displacement Load” (page 42) in Chapter “Joints and Degrees of Freedom” for more information.

Self-Weight Load

Self-Weight Load activates the self-weight of all elements in the model. Self-weight always acts downward, in the global $-Z$ direction. You may scale the self-weight by a single scale factor that applies to the whole structure. No Self-Weight Load can be produced by an element with zero weight.

For more information:

- See Topic “Upward and Horizontal Directions” (page 13) in Chapter “Coordinate Systems.”
- See Topic “Self-Weight Load” (page 134) in Chapter “The Frame Element.”
- See Topic “Self-Weight Load” (page 172) in Chapter “The Cable Element.”
- See Topic “Self-Weight Load” (page 191) in Chapter “The Shell Element.”
- See Topic “Self-Weight Load” (page 221) in Chapter “The Plane Element.”
- See Topic “Self-Weight Load” (page 232) in Chapter “The Asolid Element.”

- See Topic “Self-Weight Load” (page 248) in Chapter “The Solid Element.”
- See Topic “Self-Weight Load” (page 272) in Chapter “The Link/Support Element—Basic.”
- See Topic “Self-Weight Load” (page 317) in Chapter “The Tendon Object.”

Gravity Load

Gravity Load activates the self-weight of the Frame, Cable, Shell, Plane, Asolid, Solid, and Link/Support elements. For each element to be loaded, you may specify the gravitational multipliers **ux**, **uy**, and **uz** in any fixed coordinate system **csys**. Multiplier values are additive after being converted to the global coordinate system.

Each element produces a Gravity Load, having three components in system **csys**, equal to its self-weight multiplied by the factors **ux**, **uy**, and **uz**. This load is apportioned to each joint of the element. For example, if **uz** = -2, twice the self-weight is applied to the structure acting in the negative Z direction of system **csys**. No Gravity Load can be produced by an element with zero weight.

The difference between Self-Weight Load and Gravity Load is:

- Self-Weight Load acts equally on all elements of the structure and always in the global -Z direction
- Gravity Load may have a different magnitude and direction for each element in the structure

Both loads are proportional to the self-weight of the individual elements.

For more information:

- See Topic “Gravity Load” (page 135) in Chapter “The Frame Element.”
- See Topic “Gravity Load” (page 173) in Chapter “The Cable Element.”
- See Topic “Gravity Load” (page 207) in Chapter “The Shell Element.”
- See Topic “Gravity Load” (page 221) in Chapter “The Plane Element.”
- See Topic “Gravity Load” (page 233) in Chapter “The Asolid Element.”
- See Topic “Gravity Load” (page 249) in Chapter “The Solid Element.”
- See Topic “Gravity Load” (page 272) in Chapter “The Link/Support Element—Basic.”
- See Topic “Gravity Load” (page 318) in Chapter “The Tendon Object.”

Concentrated Span Load

Concentrated Span Load applies concentrated forces and moments at arbitrary locations on Frame elements. You may specify components **ux**, **uy**, **uz**, **rx**, **ry**, and **rz** in any fixed coordinate system **csys**, and components **u1**, **u2**, **u3**, **r1**, **r2**, and **r3** in the Frame element local coordinate system. Force values are additive after being converted to the Frame element local coordinate system.

See Topic “Concentrated Span Load” (page 135) in Chapter “The Frame Element” for more information.

Distributed Span Load

Distributed Span Load applies distributed forces and moments at arbitrary locations on Frame and Cable elements. You may specify components **ux**, **uy**, **uz**, **rx**, **ry**, and **rz** in any fixed coordinate system **csys**, and components **u1**, **u2**, **u3**, **r1**, **r2**, and **r3** in the Frame element local coordinate system. Force values are additive after being converted to the Frame element local coordinate system.

For more information, See Topic “Distributed Span Load” (page 137) in Chapter “The Frame Element”, and Topic “Distributed Span Load” (page 173) in Chapter “The Cable Element”

Tendon Prestress Load

Tendons are a special type of object that can be embedded inside other objects (frames, shells, planes, asolids, and solids) to represent the effect of prestressing and post-tensioning. These tendons attach to the other objects and impose load upon them.

You may specify whether the tendons are to be modeled as independent elements in the analysis, or just to act upon the rest of the structure as loads. This affects the types of loads that are directly imposed upon the structure.

See Topic “Prestress Load” (page 316) in Chapter “The Tendon Object” for more information.

Uniform Load

Uniform Load applies uniformly distributed forces to the mid-surface of Shell elements. You may specify components **ux**, **uy**, and **uz** in any fixed coordinate system **csys**, and components **u1**, **u2**, and **u3** in the element local coordinate system. Force values are additive after being converted to the element local coordinate system.

See Topic “Uniform Load” (page 207) in Chapter “The Shell Element” for more information.

Surface Pressure Load

Surface Pressure Load applies an external pressure to any of the outer faces of the Shell, Plane, Asolid, and Solid elements. The load on each face of an element is specified independently.

You may specify pressures, **p**, that are uniform over an element face or interpolated from pressure values given by Joint Patterns. Hydrostatic pressure fields may easily be specified using Joint Patterns. Pressure values are additive.

For more information:

- See Topic “Surface Pressure Load” (page 208) in Chapter “The Shell Element.”
- See Topic “Surface Pressure Load” (page 222) in Chapter “The Plane Element.”
- See Topic “Surface Pressure Load” (page 233) in Chapter “The Asolid Element.”
- See Topic “Surface Pressure Load” (page 249) in Chapter “The Solid Element.”
- See Topic “Joint Patterns” (page 332) in this Chapter.

Pore Pressure Load

Pore Pressure Load models the drag and buoyancy effects of a fluid *within* a solid medium, such as the effect of water upon the solid skeleton of a soil. Pore Pressure Load may be used with Shell, Asolid, and Solid elements.

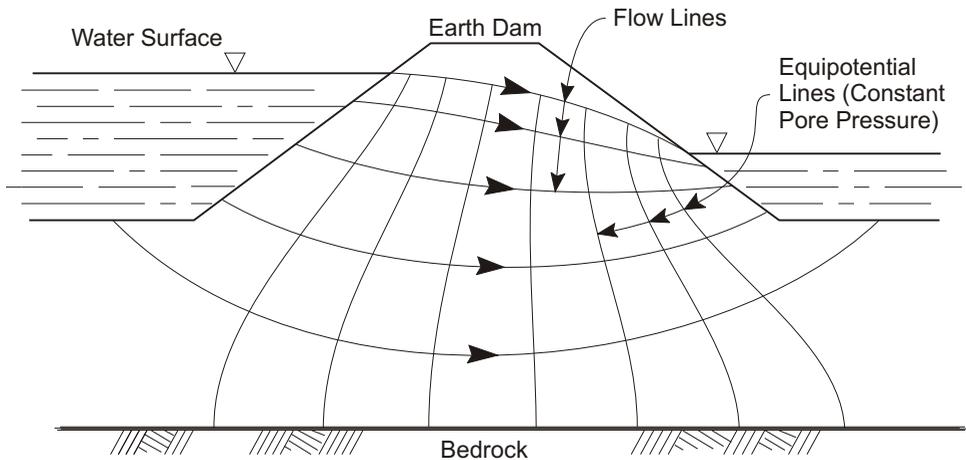


Figure 78
Flow-net Analysis of an Earth Dam to Obtain Pore Pressures

Scalar fluid-pressure values are given at the element joints by Joint Patterns, and interpolated over the element. These pressure values may typically be obtained by flow-net analysis, such as illustrated in Figure 78 (page 329). Hydrostatic pressure fields may easily be specified using Joint Patterns. Pressure values are additive.

The total force acting on the element is the integral of the gradient of this pressure field over the volume of the element. This force is apportioned to each of the joints of the element. The forces are typically directed from regions of high pressure toward regions of low pressure.

Note that although pressures are specified, it is the *pressure gradient* over an element that causes the load. Thus a uniform pressure field over an element will cause no load. Pressure differences *between* elements also cause no load. For this reason, it is important that the pore-pressure field be continuous over the structure.

The displacements, stresses, and reactions due to Pore Pressure Load represent the response of the solid medium, not that of the combined fluid and solid structure. In the case of soil, the stresses obtained are the usual “effective stresses” of soil mechanics (Terzaghi and Peck, 1967). Note, however, that the *total* soil weight and mass density should be used for the material properties.

For more information:

- See Topic “Pore Pressure Load” (page 222) in Chapter “The Plane Element.”

- See Topic “Pore Pressure Load” (page 234) in Chapter “The Asolid Element.”
- See Topic “Pore Pressure Load” (page 249) in Chapter “The Solid Element.”
- See Topic “Joint Patterns” (page 332) in this Chapter.

Temperature Load

Temperature Load creates thermal strains in the Frame, Cable, Shell, Plane, Asolid, and Solid elements. These strains are given by the product of the Material coefficients of thermal expansion and the *temperature change* of the element. All specified Temperature Loads represent a change in temperature from the unstressed state for a linear analysis, or from the previous temperature in a nonlinear analysis.

Load Temperature gradients may also be specified in the two transverse directions of the Frame element, and in the thickness direction of the Shell element. These gradients induce bending strains in the elements. Temperature gradients are specified as the linear *change* in temperature per unit length. Thus to specify a given difference in temperature across the depth of a Frame or Shell element, divide the temperature difference by the corresponding depth, and assign this value as the temperature gradient load.

The Load Temperatures and gradients may be different for each Load Pattern. You may specify temperatures and/or gradients that are uniform over an element or that are interpolated from values given by Joint Patterns.

For more information:

- See Topic “Temperature Load” (page 140) in Chapter “The Frame Element.”
- See Topic “Temperature Load” (page 174) in Chapter “The Cable Element.”
- See Topic “Temperature Load” (page 209) in Chapter “The Shell Element.”
- See Topic “Temperature Load” (page 223) in Chapter “The Plane Element.”
- See Topic “Temperature Load” (page 234) in Chapter “The Asolid Element.”
- See Topic “Temperature Load” (page 250) in Chapter “The Solid Element.”
- See Topic “Temperature Load” (page 318) in Chapter “The Tendon Object.”
- See Topic “Joint Patterns” (page 332) in this Chapter.

Strain Load

Strain Load induces distributed strains in the Frame, Cable, and Shell elements. The imposed strains tend to cause deformation in unrestrained elements, or create internal forces and stresses in restrained elements. The types of strains that are available is different for each type of element.

You may specify strains that are uniform over an element or that are interpolated from values given by Joint Patterns.

For more information:

- See Topic “Strain Load” (page 141) in Chapter “The Frame Element.”
- See Topic “Strain and Deformation Load” (page 174) in Chapter “The Cable Element.”
- See Topic “Strain Load” (page 210) in Chapter “The Frame Element.”
- See Topic “Joint Patterns” (page 332) in this Chapter.

Deformation Load

Deformation Load is an alternative form of Strain Load where the applied deformation is specified over the whole element rather than on a per-unit-length basis. Deformation Load is only available for Frame and Cable elements. The assumed distribution of strain over the element is fixed as described in the topics referenced below. Joint Patterns are not used.

For more information:

- See Topic “Deformation Load” (page 141) in Chapter “The Frame Element.”
- See Topic “Strain and Deformation Load” (page 174) in Chapter “The Cable Element.”

Target-Force Load

Target-Force Load is a special type of loading where you specify a desired axial force, and deformation load is iteratively applied to achieve the target force. Since the axial force may vary along the length of the element, you must also specify the relative location where the desired force is to occur. Target-Force loading is only

used for nonlinear static and staged-construction analysis. If applied in any other type of Load Case, it has no effect.

Unlike all other types of loading, target-force loading is not incremental. Rather, you are specifying the total force that you want to be present in the frame element at the end of the Load Case or construction stage. The applied deformation that is calculated to achieve that force may be positive, negative, or zero, depending on the force present in the element at the beginning of the analysis. When a scale factor is applied to a Load Pattern that contains Target-Force loads, the total target force is scaled. The increment of applied deformation that is required may change by a different scale factor.

For more information:

- See Topic “Target-Force Load” (page 142) in Chapter “The Frame Element.”
- See Topic “Target-Force Load” (page 174) in Chapter “The Cable Element.”
- See Topic “Target-Force Iteration” (page 444) in Chapter “Nonlinear Static Analysis” for more information.

Rotate Load

Rotate Load applies centrifugal force to Asolid elements. You may specify an angular velocity, \mathbf{r} , for each element. The centrifugal force is proportional to the square of the angular velocity. The angular velocities are additive. The load on the element is computed from the total angular velocity.

See Topic “Rotate Load” (page 234) in Chapter “The Asolid Element.”

Joint Patterns

A Joint Pattern is a named entity that consists of a set of scalar numeric values, one value for each joint of the structure. A Joint Pattern can be used to describe how pressures or temperatures vary over the structure. Joint Patterns may also be used to specify joint offsets and thickness overwrite for Shell elements.

Patterns are most effective for describing complicated spatial distributions of numeric values. Their use is optional and is not required for simple problems.

Since Pattern values are scalar quantities, they are independent of any coordinate system. The definition of a Joint Pattern by itself causes no effect on the structure.

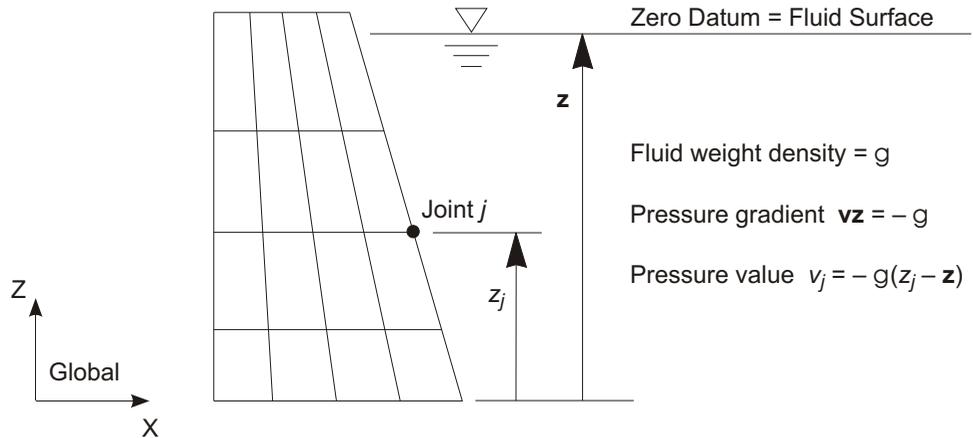


Figure 79
Example of a Hydrostatic Pressure Pattern

The pattern must be used in a pressure, temperature, or other assignment that is applied to the model.

For complicated Patterns, values could be generated in a spreadsheet program or by some other means, and brought into the model by importing tables or by using interactive table editing.

In the graphical user interface, Pattern values can be assigned to selected joints. They are specified as a linear variation in a given gradient direction from zero at a given datum point. An option is available to permit only positive or only negative values to be defined. This is useful for defining hydrostatic pressure distributions. Multiple linear variations may be assigned to the same or different joints in the structure.

The following parameters are needed for a pattern assignment:

- The components of the gradient, **A**, **B**, and **C**, in the global coordinate system
- The value **D** of the pattern at the global origin
- The choice between:
 - Setting negative values to zero
 - Setting positive values to zero
 - Allow all positive and negative values (this is the default)

The component **A** indicates, for example, how much the Pattern value changes per unit of distance parallel to the global X axis.

The Pattern value, v_j , defined for a joint j that has coordinates (x_j, y_j, z_j) is given by:

$$v_j = \mathbf{A} x_j + \mathbf{B} y_j + \mathbf{C} z_j + \mathbf{D} \quad (\text{Eqn. 1})$$

If you know the coordinates of the datum point, \mathbf{x} , \mathbf{y} , and \mathbf{z} , in global coordinate system at which the pattern value should be zero (say the free surface of water), then:

$$v_j = \mathbf{A} (x_j - \mathbf{x}) + \mathbf{B} (y_j - \mathbf{y}) + \mathbf{C} (z_j - \mathbf{z}) \quad (\text{Eqn. 2})$$

from which we can calculate that:

$$\mathbf{D} = -(\mathbf{A} \mathbf{x} + \mathbf{B} \mathbf{y} + \mathbf{C} \mathbf{z}) \quad (\text{Eqn. 3})$$

In most cases, the gradient will be parallel to one of the coordinate axes, and only one term in Eqn. 2 is needed.

For example, consider a hydrostatic pressure distribution caused by water impounded behind a dam as shown in Figure 79 (page 333). The Z direction is up in the global coordinate system. The pressure gradient is simply given by the fluid weight density acting in the downward direction. Therefore, $\mathbf{A} = 0$, $\mathbf{B} = 0$, and $\mathbf{C} = -62.4 \text{ lb/ft}^3$ or -9810 N/m^3 .

The zero-pressure datum can be any point on the free surface of the water. Thus \mathbf{z} should be set to the elevation of the free surface in feet or meters, as appropriate, and $\mathbf{D} = -\mathbf{C} \mathbf{z}$. For hydrostatic pressure, you would specify that negative values be ignored, so that any joints above the free surface will be assigned a zero value for pressure.

Mass Source

The mass used as inertia in dynamic analyses or to calculate certain types of loads can be controlled by specifying the Mass Source. Multiple Mass Sources can be defined so that different load cases can use a different mass distributions for loading and inertia. Examples could include modeling a structure supporting different configurations of equipment, or explicitly considering the effect of different eccentricities of the story mass on the mode shapes. However, for most models only a single Mass Source is needed, and this is the default.

There are three possible contributions to a Mass Source:

- **Element Self Mass** — This includes the mass from the section properties used by the elements. For the Link/Support elements, this mass is explicitly defined in the section property. For all other elements the mass comes from the material property referenced by the section property.
- **Additional Mass** — This includes mass assigned to the joints and any additional mass assigned to the frame or shell elements.
- **Specified Load Patterns** — Mass is computed from the gravity load as computed from a specified linear combination of load patterns.

For each defined Mass Source you may specify which of the above contributions are to be considered. By default, there is only a single Mass Source, and this includes the Element Self Mass and the Additional Mass, with no contribution from Specified Load Patterns.

Mass from Specified Load Patterns

When defining a Mass Source, you may include the mass from a linear combination of Load Patterns. For each Load Pattern in this combination you specify a scale factor that multiplies the loads. The mass is calculated as follows:

- For each individual joint or element, joint loads are calculated based on the specified linear combination of load patterns.
- For each joint of this element, the combined load acting in the gravity direction (negative global Z) is divided by the gravitational constant, and then assigned equally to the translational degrees of freedom at that joint.
- Downward loads ($-Z$) create positive mass, and upward loads ($+Z$) create negative mass.

Important notes:

- Loads are calculated for each element in its original configuration, independent of any large-deflection analysis.
- The mass calculated from loads acts in all three translational directions even though the load creating the mass acts vertically.
- Self-equilibrating loads, such as temperature, strain, and prestressing, should not be used for Mass Source.
- *Be careful not to double-count the self-mass by specifying both Element Self mass and a load pattern that contains self-weight.*

Negative Mass

All mass values used for analysis must be positive or zero. Negative mass is not realistic and is not allowed. However, you can use negative mass in the mass source if your intention is to subtract it from other positive mass.

The contributions due to Element Self Mass and Additional Mass come from materials, link properties, joint mass, line mass, and area mass. Except for joint mass, these contributions are always positive and are equal for the three translational degrees of freedom at each joint. Joint mass can be positive or negative and may be different at each of the six joint degrees of freedom.

The contributions from Specified Load Patterns are equal for the three translational degrees of freedom at each joint. Downward loads generate positive mass and upward loads generate negative mass.

The net result after combining all these contributions for any joint degree of freedom may be positive or negative and will be reported as such in the “Assembled Joint Mass” table. During equation solution, any negative mass values that remain after combining the mass for constraints are set to zero and a warning issued in the analysis log file.

Multiple Mass Sources

For most models, a single Mass Source is sufficient and is defined as described above. This will apply to all loads calculated from mass and to all dynamic load cases.

However, you have the option (when available) to specify multiple Mass Sources for use with different load cases. Each Mass Source is defined as described above by specifying the three contributions to the mass. One Mass Source is selected as the default and this will be used for all Load Cases unless specified otherwise.

A specified Mass Source can be selected for the following types of load cases:

- Nonlinear static
- Nonlinear staged-construction
- Nonlinear direct-integration time-history

If you do not select a specified Mass Source for these types of load cases, the Mass source used will be the default Mass Source if the load case starts from zero initial conditions, or the Mass Source used in the previous load case if continuing from another one of these same three types of load cases.

For all other load cases, a similar rule applies. The default Mass Source will be used if the load case uses the stiffness from zero load conditions. Otherwise, the Mass Source used will be the same as for the nonlinear static, staged-construction, or direct-integration time-history load case used to calculate the stiffness. Response-spectrum and modal time-history load cases use the Mass Source of their corresponding modal load case. See Chapter “Load Cases” (page 341) for more information, especially Topic “Sequence of Analysis” (page 344).

For example, consider the case where a response-spectrum analysis is to be carried out for a tower both with and without a significant equipment load. You could do the following:

- Define two load patterns
 - DEAD, which includes the self-weight of the tower structure
 - LIVE, which includes only the weight of the equipment
- Define two mass sources
 - MASSDEAD, which includes only the load pattern DEAD with a scale factor of 1.0, and no Element Self Mass or Additional Mass
 - MASSDEADLIVE, which includes both load patterns DEAD and LIVE, each with a scale factor of 1.0, and no Element Self Mass or Additional Mass
- Define two nonlinear static load cases
 - DEAD, which specifies mass source MASSDEAD
 - DEADLIVE, which specifies mass source MASSDEADLIVE
- Define two modal load cases
 - MODALDEAD, which uses the stiffness of load case DEAD
 - MODALDEADLIVE, which uses the stiffness of load case DEADLIVE
- Define two response-spectrum load cases
 - SPECDEAD, which uses the modes of load case MODALDEAD
 - DEADLIVE, which uses the modes of load case MODALDEADLIVE

Note that in the above example the nonlinear static load cases were used only to specify the Mass Source. However, in most practical cases you would also want to apply the corresponding load patterns as loads and consider P-delta effects, as these would likely also have an effect on the modes.

Note that the response-spectrum cases, in addition to considering the inertial effect of the different masses on the modes, also apply acceleration loads that are based on

mass. These loads will automatically be based on the mass from the same Mass Source used to calculate the modes. See topic “Acceleration Load” (page 338).

Automated Lateral Loads

Automated lateral loads are available as specified in the separate “Automated Lateral Loads Manual” provided with the software. Most of the seismic loads, and some wind loads, are calculated based on mass, which in turn is calculated by analysis from the Mass Source.

When multiple Mass Sources are used, the mass could be different between load cases. For this reason, any given load pattern that contains an automated lateral load based on mass can only be applied in load cases that use the same mass source. If the same load pattern is applied in multiple load cases using different mass sources, a warning will be issued when the analysis is run and only those load cases using one of the mass sources will be run. In this case, you can define duplicate load patterns as needed for use with load cases using different mass sources.

Some automated lateral loads use a modal period that can be specified by the user or determined from analysis. By default, this period will be determined from the first modal load case that is run, whether or not it uses the same Mass Source as other load cases that apply the load pattern. For this reason, when using multiple Mass Sources, it is strongly recommended that you review the modal period used for automated lateral loads and specify your own value to be used if the calculated value does not apply.

Acceleration Loads

In addition to the Load Patterns that you define, the program automatically computes three Acceleration Loads that act on the structure due to unit translational accelerations in each of the global directions, and three unit rotational accelerations about the global axes at the global origin. Acceleration Loads can be applied in a Load Case just like Load Patterns.

Acceleration Loads are determined by d’Alembert’s principal, and are denoted m_x , m_y , m_z , m_{rx} , m_{ry} , and m_{rz} . These loads are used for applying ground accelerations in response-spectrum (translation only) and time-history analyses, and can be used as starting load vectors for Ritz-vector analysis.

These loads are computed for each joint and element and summed over the whole structure. The translational Acceleration Loads for the joints are simply equal to the

negative of the joint translational masses in the joint local coordinate system. These loads are transformed to the global coordinate system.

The translational Acceleration Loads for all elements except for the Asolid are the same in each direction and are equal to the negative of the element mass. No coordinate transformations are necessary. Rotational acceleration will generally differ about each axis.

For the Asolid element, the Acceleration Load in the global direction corresponding to the axial direction is equal to the negative of the element mass. The Acceleration Loads in the radial and circumferential directions are zero, since translations in the corresponding global directions are not axisymmetric. Similar considerations apply to the rotational accelerations.

The Acceleration Loads can be transformed into any coordinate system. In a fixed coordinate system (global or Alternate), the translational Acceleration Loads along the positive X, Y, and Z axes are denoted UX, UY, and UZ, respectively; the rotational Acceleration Loads about the X, Y, and Z axes are similarly denoted RX, RY, and RZ.

In a local coordinate system defined for a response-spectrum or time-history analysis, the Acceleration Loads along or about the positive local 1, 2, and 3 axes are denoted U1, U2, U3, R1, R2, and R3 respectively. Rotational accelerations will be applied about the origin of the coordinate system specified with the Acceleration Load in the Load Case. Each Acceleration Load applied in a given Load Case can use a separate coordinate system.

Although the rotational accelerations are applied in load cases about the origins of the specified coordinate systems, the modal participation factors and the various modal participation ratios for the accelerations RX, RY, and RZ are reported about the center of mass for the structure:

- For the RX load, this location is the y centroid of the UZ load and the z centroid of the UY load.
- For the RY load, this location is the z centroid of the UX load and the x centroid of the UZ load.
- For the RZ load, this location is the z centroid of the UY load and the y centroid of the UX load.

This makes the rotational participation factors and ratios are more meaningful because they do not include any contribution from the translational accelerations.

Chapter XVIII

Load Cases

A Load Case defines how the loads are to be applied to the structure (e.g., statically or dynamically), how the structure responds (e.g., linearly or nonlinearly), and how the analysis is to be performed (e.g., modally or by direct-integration.)

Basic Topics for All Users

- Overview
- Load Cases
- Types of Analysis
- Sequence of Analysis
- Running Load Cases
- Linear and Nonlinear Load Cases
- Linear Static Analysis
- Functions
- Load Combinations (Combos)
- Global Energy Response

Advanced Topics

- Multi-Step Static Analysis

- Linear Buckling Analysis
- Equation Solvers
- Environment Variables to Control Analysis
- Accessing the Assembled Stiffness and Mass Matrices

Overview

A Load Case defines how loads are to be applied to the structure, and how the structural response is to be calculated. You may define as many named Load Cases of any type that you wish. When you analyze the model, you may select which cases are to be run. You may also selectively delete results for any Load Case.

Note: Load Patterns by themselves do not create any response (deflections, stresses, etc.) You must define a Load Case to apply the load.

There are many different types of Load Cases. Most broadly, analyses are classified as linear or nonlinear, depending upon how the structure responds to the loading.

The results of linear analyses may be superposed, i.e., added together after analysis. The available types of linear analysis are:

- Static analysis
- Modal analysis for vibration modes, using eigenvectors or Ritz vectors
- Response-spectrum analysis for seismic response
- Time-history dynamic response analysis
- Buckling-mode analysis
- Moving-load analysis for bridge and other vehicle live loads
- Steady-state analysis
- Power-spectral-density analysis

The results of nonlinear analyses should not normally be superposed. Instead, all loads acting together on the structure should be combined directly within the Load Cases. Nonlinear Load Cases may be chained together to represent complex loading sequences. The available types of nonlinear analyses are:

- Nonlinear static analysis
- Nonlinear time-history analysis

Named Load Combinations can also be defined to combine the results of Load Cases. Results can be combined additively or by enveloping. Additive Load Combinations of nonlinear Load Cases is not usually justified.

Load Cases

Each different analysis performed is called a Load Case. For each Load Case you define, you need to supply the following type of information:

- **Case name:** This name must be unique across all Load Cases of all types. The case name is used to request analysis results (displacements, stresses, etc.), for creating Load Combinations, and sometimes for use by other dependent Load Cases.
- **Analysis type:** This indicate the type of analysis (static, response-spectrum, buckling, etc.), as well as available options for that type (linear, nonlinear, etc.).
- **Prerequisite load cases:** Some load cases may continue from a previous load case, use the stiffness from a previous load case, and/or use the modes from a previous load case.
- **Loads applied:** For most types of analysis, you specify the Load Patterns that are to be applied to the structure.

Additional data may be required, depending upon the type of analysis being defined.

Types of Analysis

There are many different types of Load Cases. Most broadly, analyses are classified as linear or nonlinear, depending upon how the structure responds to the loading. See Topic “Linear and Nonlinear Load Cases” (page 346) in this Chapter.

The results of linear analyses may be superposed, i.e., added together after analysis. The available types of linear analysis are:

- Static analysis
- Modal analysis
- Response-spectrum analysis
- Time-history analysis, by modal superposition or direct integration

- Buckling analysis
- Moving-load analysis
- Steady-state analysis
- Power-spectral-density analysis

The results of nonlinear analyses should not normally be superposed. Instead, all loads acting together on the structure should be combined directly within the Load Cases. Nonlinear Load Cases may be chained together to represent complex loading sequences. The available types of nonlinear analyses are:

- Nonlinear static analysis
- Nonlinear time-history analysis, by modal superposition or direct integration

After you have defined a Load Case, you may change its type at any time. When you do, the program will try to carry over as many parameters as possible from the old type to the new type. Parameters that cannot be carried over will be set to default values, which you can change.

For more information:

- See Topic “Linear Static Analysis” (page 348) in this Chapter
- See Topic “Buckling Analysis” (page 349) in this Chapter
- See Chapter “Modal Analysis” (page 365)
- See Chapter “Response-Spectrum Analysis” (page 383)
- See Chapter “Linear Time-History Analysis” (page 397)
- See Chapter “Nonlinear Static Analysis” (page 425)
- See Chapter “Nonlinear Time-History Analysis” (page 447)
- See Chapter “Frequency-Domain Analysis” (page 465)
- See Chapter “Moving-Load Analysis” (page 477)

Sequence of Analysis

A Load Case may be dependent upon other Load Cases in the following situations:

- A modal-superposition type of Load Case (response-spectrum or modal time-history) uses the modes calculated in a modal Load Case
- A nonlinear Load Case may continue from the state at the end of another nonlinear case

- A linear Load Cases may use the stiffness of the structure as computed at the end of a nonlinear case

A Load Case which depends upon another case is called **dependent**. The case upon which it depends is called a **prerequisite** case.

When the program performs analysis, it will always run the cases in the proper order so that dependent cases are run after any of their prerequisite cases.

You can build up one or more sequences of Load Cases that can be as simple or as complicated as you need. However, each sequence must ultimately start with a Load Case that itself starts from zero and does not have any prerequisite cases.

Example

A common example would be to define a nonlinear static Load Case with the following main features:

- The name is, say, “PDELTA”
- The type is nonlinear static
- The loads applied are Load Pattern “DEAD” scaled by 1.0 plus Load Pattern “LIVE” scaled by 0.25. These represent a typical amount of gravity load on the structure
- The only nonlinearity considered is the P-delta effect of the loading

We are not necessarily interested in the response of Load Case PDELTA, but rather we will use the stiffness at the end this case for a series of linear Load Cases that we are interested in. These may include linear static cases for all loads of interest (dead, live, wind, snow, etc.), a modal Load Case, and a response-spectrum Load Case.

Because these cases have all been computed using the same stiffness, their results are superposable, making it very simple to create any number of Load Combinations for design purposes.

Running Load Cases

After you have defined a structural model and one or more Load Cases, you must explicitly **run** the Load Cases to get results for display, output, and design purposes.

When an analysis is run, the program converts the object-based model to finite elements, and performs all calculations necessary to determine the response of the

structure to the loads applied in the Load Cases. The analysis results are saved for each case for subsequent use.

By default, all Load Cases defined in the model are run each time you perform an analysis. However, you can change this behavior. For each Load Case, you can set a flag that indicates whether or not it will be run the next time you perform an analysis. This enables you to define as many different cases as you need without having to run all of them every time. This is particularly useful if you have nonlinear cases that may take a long time to run.

If you select a case to be run, the program will also run all prerequisite cases that have not yet been run, whether you select them or not.

You can create new Load Cases without deleting the results of other cases that have already been run. You can also modify existing Load Cases. However, the results for the modified case and all cases that depend upon it will be deleted.

When an analysis is performed, the cases will be run in an order that is automatically determined by the program in order to make sure all prerequisite cases are run before their dependent cases. If any prerequisite cases fail to complete, their dependent cases will not be run. However, the program will continue to run other cases that are not dependent upon the failed cases.

You should always check the analysis log (.LOG) file to see statistics, warnings, and error messages that were reported during the analysis. You can also see a summary of the cases that have been run, and whether or not they completed successfully, using the Analysis commands in the graphical user interface.

Whenever possible, the program will re-use the previously solved stiffness matrix while running Load Cases. Because of this, the order in which the cases are run may not be the same each time you perform an analysis.

See Topic “Sequence of Analysis” (page 467) in this Chapter for more information.

Linear and Nonlinear Load Cases

Every Load Case is considered to be either linear or nonlinear. The difference between these two options is very significant in SAP2000, as described below.

All Load Case types may be linear. Only static analysis and time-history analysis may be nonlinear.

Structural properties

Linear: Structural properties (stiffness, damping, etc.) are constant during the analysis.

Nonlinear: Structural properties may vary with time, deformation, and loading. How much nonlinearity actually occurs depends upon the properties you have defined, the magnitude of the loading, and the parameters you have specified for the analysis.

Initial conditions

Linear: The analysis starts with zero stress. It does not contain loads from any previous analysis, even if it uses the stiffness from a previous nonlinear analysis.

Nonlinear: The analysis may continue from a previous nonlinear analysis, in which case it contains all loads, deformations, stresses, etc., from that previous case.

Structural response and superposition

Linear: All displacements, stresses, reactions, etc., are directly proportional to the magnitude of the applied loads. The results of different linear analyses may be superposed.

Nonlinear: Because the structural properties may vary, and because of the possibility of non-zero initial conditions, the response may not be proportional to the loading. Therefore, the results of different nonlinear analyses should not usually be superposed.

Linear Static Analysis

The linear static analysis of a structure involves the solution of the system of linear equations represented by:

$$K u = r$$

where K is the stiffness matrix, r is the vector of applied loads, and u is the vector of resulting displacements. See Bathe and Wilson (1976).

You may create any number of linear static Load Cases. For each case you may specify a combination of one or more Load Patterns and/or Acceleration Loads that

make up the load vector r . Normally, you would specify a single Load Pattern with a scale factor of one.

Every time you define a new Load Pattern, the program automatically creates a corresponding linear static Load Case of the same name. This Load Case applies the Load Case with a unit scale factor. If you delete or modify the Load Case, the analysis results will not be available, even though the Load Case may still exist.

For a new model, the program creates a default Load Pattern call DEAD which applies the self weight of the structure. The corresponding linear static Load Case is also called DEAD.

For each linear static Load Case, you may specify that the program use the stiffness matrix of the full structure in its unstressed state (the default), or the stiffness of the structure at the end of a nonlinear Load Case. The most common reasons for using the stiffness at the end of a nonlinear case are:

- To include P-delta effects from an initial P-delta analysis
- To include tension-stiffening effects in a cable structure
- To consider a partial model that results from staged construction

See Chapter “Nonlinear Static Analysis” (page 425) for more information.

Multi-Step Static Analysis

Certain types of Load Patterns are multi-stepped, meaning that they actually represent many separate spatial loading patterns applied in sequence. These include the Vehicle-Live and Wave types of Load Patterns.

You can apply multi-stepped Load Patterns in a Multi-Step Static Load Case. This type of Load Case will perform a series of independent linear analyses, one for each step of the Load Pattern:

$$K u_i = r_i$$

where r_i is the load at step I , and u_i is the resulting solution.

You can apply any linear combination of Load Patterns in the same Load Case, each with an independent scale factor. These will be summed together as follows: all multi-stepped Load Patterns will be synchronized, applying the load for the same step number at the same time, and all single-stepped Load Patterns (e.g., dead load) will be applied in every step.

The number of solution steps for the Load Case will be equal to the largest number of load steps of any of the applied multi-stepped Load Patterns.

Although the multi-stepped Load Patterns represent a time series of loads, Multi-Step Static Analysis does not include any dynamical effects. You can consider dynamics by converting the Multi-Step Static Load Case to a Time-History Load Case, which can be linear or nonlinear.

Linear Buckling Analysis

Linear buckling analysis seeks the instability modes of a structure due to the P-delta effect under a specified set of loads. Buckling analysis involves the solution of the generalized eigenvalue problem:

$$[\mathbf{K} \quad \mathbf{G}(\mathbf{r})] \quad \mathbf{0}$$

where \mathbf{K} is the stiffness matrix, $\mathbf{G}(\mathbf{r})$ is the geometric (P-delta) stiffness due to the load vector \mathbf{r} , λ is the diagonal matrix of eigenvalues, and \mathbf{u} is the matrix of corresponding eigenvectors (mode shapes).

Each eigenvalue-eigenvector pair is called a buckling mode of the structure. The modes are identified by numbers from 1 to n in the order in which the modes are found by the program.

The eigenvalue λ is called the buckling factor. It is the scale factor that must multiply the loads in \mathbf{r} to cause buckling in the given mode. It can also be viewed as a safety factor: if the buckling factor is greater than one, the given loads must be increased to cause buckling; if it is less than one, the loads must be decreased to prevent buckling. The buckling factor can also be negative. This indicates that buckling will occur if the loads are reversed.

You may create any number of linear buckling Load Cases. For each case you specify a combination of one or more Load Patterns and/or Acceleration Loads that make up the load vector \mathbf{r} . You may also specify the number of modes to be found and a convergence tolerance. It is strongly recommended that you seek more than one buckling mode, since often the first few buckling modes may have very similar buckling factors. A minimum of six modes is recommended.

It is important to understand that buckling modes depend upon the load. There is not one set of buckling modes for the structure in the same way that there is for natural vibration modes. You must explicitly evaluate buckling for each set of loads of concern.

For each linear buckling Load Case, you may specify that the program use the stiffness matrix of the full structure in its unstressed state (the default), or the stiffness of the structure at the end of a nonlinear Load Case. The most common reasons for using the stiffness at the end of a nonlinear case are:

- To include P-delta effects from an initial P-delta analysis
- To include tension-stiffening effects in a cable structure
- To consider a partial model that results from staged construction

See Chapter “Nonlinear Static Analysis” (page 425) for more information.

Functions

A Function is a series of digitized abscissa-ordinate pairs that may represent:

- Pseudo-spectral acceleration vs. period for response-spectrum analysis, or
- Load vs. time for time-history analysis
- Load vs. frequency for steady-state analysis
- Power density (load squared per frequency) vs. frequency for power-spectral-density analysis

You may define any number of Functions, assigning each one a unique label. You may scale the abscissa and/or ordinate values whenever the Function is used.

The abscissa of a Function is always time, period, or frequency. The abscissa-ordinate pairs must be specified in order of increasing abscissa value.

If the increment between abscissa values is constant and the Function starts at abscissa zero, you need only specify the abscissa increment, **dt**, and the successive function values (ordinates) starting at abscissa zero. The function values are specified as:

f0, f1, f2, ..., fn

corresponding to abscissas:

0, dt, 2 dt, ..., n dt

where $n + 1$ is the number of values given.

If the abscissa increment is not constant or the Function does not start at abscissa zero, you must specify the pairs of abscissa and function value as:

$t_0, f_0, t_1, f_1, t_2, f_2, \dots, t_n, f_n$

where $n + 1$ is the number of pairs of values given.

Load Combinations (Combos)

A Load Combination (Combo) is a named combination of the results from Load Cases or other Load Combinations. Combo results include all displacements and forces at the joints and internal forces or stresses in the elements.

You may define any number of Combos. To each one of these you assign a unique name, that also should not be the same as any Load Case. Combos can combine the results of Load Cases and also those of other Combos, provided a circular dependency is not created.

By default, each Combo produces a pair of values for each response quantity: a maximum and a minimum. These two values may be equal for certain type of Combos, as discussed below.

Alternatively, correspondence may be requested for tabular output, in which case multiple pairs of maximum and minimum values may be produced at each output location for each Combo. This is discussed in more detail in the subtopic “Correspondence” below. The discussion that immediately follows next considers the default case where a single maximum and minimum value is produced for each response quantity.

Contributing Cases

Each contributing Load Case or Combo supplies one or two values to the Combo for each response quantity:

- Linear static cases, individual modes from Modal or Buckling cases, individual steps from multi-stepped Load Cases, and additive Combos of these types of results provide a single value. For the purposes of defining the Combos below, this single value can be considered to be two equal values
- Response-spectrum cases provide two values: the maximum value used is the positive computed value, and the minimum value is just the negative of the maximum.
- Envelopes of results from multi-stepped Load Cases provide two values: a maximum and minimum.

- For Moving-Load Load Cases, the values used are the maximum and minimum values obtained for any vehicle loading of the lanes permitted by the parameters of the analysis.

For some types of Combos, both values are used. For other types of Combos, only the value with the larger absolute value is used.

Each contributing case is multiplied by a scale factor, **sf**, before being included in a Combo.

Types of Combos

Five types of Combos are available. For each individual response quantity (force, stress, or displacement component) the two Combo values are calculated as follows:

- Additive type: The Combo maximum is an algebraic linear combination of the maximum values for each of the contributing cases. Similarly, Combo minimum is an algebraic linear combination of the minimum values for each of the contributing cases.
- Absolute type: The Combo maximum is the sum of the larger absolute values for each of the contributing cases. The Combo minimum is the negative of the Combo maximum.
- SRSS type: The Combo maximum is the square root of the sum of the squares of the larger absolute values for each of the contributing cases. The Combo minimum is the negative of the Combo maximum.
- Range type: The Combo maximum is the sum of the positive maximum values for each of the contributing cases (a case with a negative maximum does not contribute.) Similarly, the Combo minimum is the sum of the negative minimum values for each of the contributing cases (a case with a positive minimum does not contribute.)
- Envelope type: The Combo maximum is the maximum of all of the maximum values for each of the contributing cases. Similarly, Combo minimum is the minimum of all of the minimum values for each of the contributing cases.

Only additive Combos of single-valued Load Cases produce a single-valued result, i.e., the maximum and minimum values are equal. All other Combos will generally have different maximum and minimum values.

Examples

For example, suppose that the values, after scaling, for the displacement at a particular joint are 3.5 for Linear Static Load Case LL and are 2.0 for Response-spectrum case QUAKE. Suppose that these two cases have been included in an additive-type Combo called COMB1 and an envelope-type Combo called COMB2. The results for the displacement at the joint are computed as follows:

- COMB1: The maximum is $3.5 + 2.0 = 5.5$, and the minimum is $3.5 - 2.0 = 1.5$
- COMB2: The maximum is $\max(3.5, 2.0) = 3.5$, and the minimum is $\min(3.5, -2.0) = -2.0$

As another example, suppose that Linear Static Cases GRAV, WINDX and WINDY are gravity load and two perpendicular, transverse wind loads, respectively; and that a response-spectrum case named EQ has been performed. The following four Combos could be defined:

- WIND: An SRSS-type Combo of the two wind loads, WINDX and WINDY. The maximum and minimum results produced for each response quantity are equal and opposite
- GRAVEQ: An additive-type Combo of the gravity load, GRAV, and the response-spectrum results, EQ. The Combo automatically accounts for the positive and negative senses of the earthquake load
- GRAVWIN: An additive-type Combo of the gravity load, GRAV, and the wind load given by Combo WIND, which already accounts for the positive and negative senses of the load
- SEVERE: An envelope-type Combo that produces the worst case of the two additive Combos, GRAVEQ and GRAVWIN

Suppose that the values of axial force in a frame element, after scaling, are 10, 5, 3, and 7 for cases GRAV, WINDX, WINDY, and EQ, respectively. The following results for axial force are obtained for the Combos above:

- WIND: maximum $\sqrt{5^2 + 3^2} = 5.8$, minimum -5.8
- GRAVEQ: maximum $10 + 7 = 17$, minimum $10 - 7 = 3$
- GRAVWIN: maximum $10 + 5.8 = 15.8$, minimum $10 - 5.8 = 4.2$
- SEVERE: maximum $\max(17, 15.8) = 17$, minimum $\min(3, 4.2) = 3$

Range-type Combos enable you to perform skip-pattern loading very efficiently. For example, suppose you have a four-span continuous beam, and you want to

know what pattern of uniform loading on the various spans creates the maximum response:

- Create four Load Patterns, each with uniform loading on a single span
- Create four corresponding Linear Static Load Cases, each applying a single Load Pattern.
- Create a single range Combo, combining the results of the four Load Cases

The effect of this Combo is the same as enveloping all additive combinations of any single span loaded, any two spans loaded, any three spans loaded, and all four spans loaded. This range Combo could be added or enveloped with other Load Cases and Combos.

As you can see, using Combos of Combos gives you considerable power and flexibility in how you can combine the results of the various Load Cases.

Correspondence

When correspondence is requested for tabular output, the maximum and minimum value is produced for each response quantity as described above, but these are presented with the values of other related response quantities at the same location that occur at the same time the given maximum or minimum is achieved.

For example, when considering Frame-element forces, the maximum value of axial force P at a given output station is presented with the corresponding values of V2, V3, T, M2, and M3 at the same output station. Six corresponding values are also presented for the minimum value of P, and so on for the maximum and minimum values of V2, V3, T, M2, and M3.

Without correspondence, the following tabular results would be presented for a given output station in a given Combo:

Type	P	V2	V3	T	M2	M3
Max	100	44	5	-6	0	1245
Min	-50	22	-5	-9	-123	-333

With correspondence, the tabular results for the same output station location would be presented as:

Type	P	V2	V3	T	M2	M3
PMax	100	32	-4	-6	-123	-111
PMin	-50	24	3	-6	0	999
V2Max	81	44	5	-7	-65	328
V2Min	-30	22	-5	-6	-44	172
V3Max	81	44	5	-7	-65	328
V3Min	-30	22	-5	-6	-44	172
TMax	50	40	0	-6	-69	1033
TMin	87	25	-4	-9	-77	-197
M2Max	-20	39	0	-8	0	1158
M2Min	27	29	-4	-7	-123	890
M3Max	71	36	5	-6	0	1245
M3Min	-38	31	-5	-9	-123	-333

In this example, row PMax presents the maximum value of axial force P and the corresponding values of the other frame forces and moments that occur at the same time as maximum P. For a Combo that is the envelope of several linear static load cases, this would be the frame forces and moments from the contributing load case that had the maximum value of P. For an additive Combo of a dead load case plus a time-history load case, this would be the summed frame forces and moments from that time step in the time-history case that, when combined with dead load, produces that maximum value of P. Similarly, row PMin is the minimum value of P and its corresponding forces and moments, and so on for V2Max through M3Min.

The shaded values in the table above indicate the maximum and minimum values that would be presented in the table without correspondence. All other values in a given column will be within the range of these two values.

It is possible, in some cases, that the maximum or minimum value of a response quantity could occur at more than one instance in a Combo, in which case the corre-

sponding values will be presented for the first instance encountered. In this case, the corresponding results may not be unique.

Correspondence is available for tabular output of the following quantities in software products that offer this option:

- Joint displacements
- Joint reactions
- Frame forces, Shell forces
- Shell stresses, Plane stresses, Asolid stresses, Solid stresses
- Link/Support forces and deformations
- Element joint forces
- Section Cuts
- Base reactions
- Bridge object section forces, Bridge object girder forces (CSiBridge only)

For each of the above, the correspondence is presented between the various components of the same type (e.g., 6 displacements, 6 forces, or 6 stresses) at a given location (joint or output station). Correspondence is not presented between components at different locations or between components of different types (e.g., displacement and force). The only exception to this is the Link/Support element, where correspondence is presented between the 6 deformation and 12 force components of a given element. If you want to see the full corresponding state of the structure, you must use step-by-step analysis.

Correspondence is available for Combos of types Additive, Envelope, and Range, provided that the contributing load cases and Combos also support correspondence.

Correspondence is not available for Combos of types Absolute and SRSS, nor for Combos that contain any of the following types of contributing load cases and/or Combos:

- Response-spectrum load case
- Steady-state load case
- PSD load case
- Moving-load case without correspondence
- Absolute Combo
- SRSS Combo

- A Combo that contains any of the above
- In addition, correspondence is not available for base reactions or element joint forces for any Combo that contains a moving-load case.

Requesting correspondence for tabular output does not affect any other type of display, design, or the Open API functions. Correspondence for design purposes is controlled by the Design Preferences. See the various design manuals that are included with the software for more information

Additional Considerations

Moving-Load Load Cases should not normally be added together, in order to avoid multiple loading of the lanes. Additive combinations of Moving Loads should only be defined within the Moving-Load Load Case itself. Therefore, it is recommended that only a single Moving Load be included in any additive-, absolute-, SRSS-, or range-type Combo, whether referenced directly as a Moving Load or indirectly through another Combo. Multiple Moving Loads may be included in any envelope-type Combo, since they are not added.

Nonlinear Load Cases should not normally be added together, since nonlinear results are not usually superposable. Instead, you should combine the applied loads within a nonlinear Load Case so that their combined effect can be properly analyzed. This may require defining many different Load Cases instead of many different Combos. Nonlinear Load Cases may be included in an envelope-type Combo, since they are not added.

Global Energy Response

Energy response is provided for the entire structure for certain types of load cases, including linear and nonlinear modal time-history, nonlinear static and staged construction, and nonlinear direct-integration time history (but see Limitations below). Energy response can be plotted and produced in output tables.

Three basic types of energy are considered: the work done by external agents that put energy into the structure, conservative energy in the structure that can be recovered, and unconservative energy that is dissipated by the structure. Looking at the energy response can tell you a lot about the behavior of the structure, but it is most useful for determining the effectiveness of energy-dissipating mechanisms that are intentionally provided as part of the design. This can include dampers, isolators, and plastic hinging in structural members.

Global Energy Response Components

The following types of energy are reported for each considered load case:

- **Input.** This is the work done by external agents that put energy into the structure. This includes:
 - Forces, moments, pressures, self-weight, and similar loads
 - Ground acceleration
 - Ground displacement
 - Element strains and displacements, including thermal and time-dependent strains (but see Limitations below).
 - P-delta effects, such as the work done by vertical forces through lateral displacements when P-delta nonlinearity is considered
- **Potential.** This is a conservative type of energy in the structure due to elastic strain that can be recovered when the loading is removed. It is often desirable to limit potential energy to the reasonable extent possible, since it represents forces and stresses to be designed against.
- **Kinetic.** This is a conservative type of energy due to the motion of the structure. It is often desirable to limit kinetic energy since this can represent the potential for damage to structural and non-structural items. Kinetic energy is always zero in a static analysis.
- **Nonlinear Viscous Damping.** This represents the viscous energy dissipated by the velocity-based damper elements. Nonlinear viscous damping does not change during nonlinear static and staged construction load cases. The following types of link elements can generate this energy component during time-history load cases:
 - Linear damping properties
 - Damper – Exponential
 - Damper – Bilinear
 - The three Friction-Pendulum Isolators (axial degree of freedom only)
- **Nonlinear Hysteretic Damping.** This represents the hysteretic energy dissipated by the displacement-based damper elements and the plastic deformation of materials (including hinges). Nonlinear hysteretic damping can be generated by the following elements during nonlinear static, staged-construction, and time-history load cases:
 - Link elements having the following types of properties:
 - Damper – Friction Spring

Multilinear Plastic

Plastic (Wen)

Rubber Isolator

The three Friction-Pendulum Isolators (shear degrees of freedom)

- Frame nonlinear hinges of all types
- Layered shell elements with nonlinear material properties
- **Global Damping.** This represents the energy dissipated by the structure as a whole. For modal time-history load cases, this includes the modal damping specified for the load case as well as modal damping specified for the materials, if any. For nonlinear direct-integration time-history load cases, this includes stiffness- and mass-proportional damping specified for the load case as well as such damping specified for the materials, if any. Global damping is viscous (velocity-based) and does not change during nonlinear static and staged-construction load cases.
- **Error.** It is expected that the sum of the potential, kinetic, and three damping energies be equal to the input energy. The energy error is calculated as the difference between the input energy and this sum. There are several reasons why this error may not be zero, such as:
 - Discretization error due to larger steps sizes for P-delta or large-displacement geometric nonlinearity. This can be checked by re-running the load case with a smaller step size.
 - Round-off error in numerically sensitive or ill-conditioned systems.
 - Energy that is excluded due to the Limitations listed below, such as for certain types of strain loading.

Remarks

Following are some general remarks considering the nature and use of these global energy response components:

- The input energy can increase or decrease as the sum of the conservative energy increases or decreases. However, the three damping energies represent work done that is not recoverable, and these can only increase the input energy.
- For purposes of design, it is generally preferable that the primary energy dissipation be in the form of nonlinear viscous and hysteretic damping, meaning that it comes from devices (dampers, isolators, etc.) and hinging mechanisms. Global damping represents other effects that are not explicitly modeled, such as from non-structural elements and the intrinsic nature of the materials. If this

is too large, then it is difficult to know whether or not the design intent is being met.

- When an object is removed during staged construction, its potential energy is removed and is also deducted from the input energy. Any damping energy produced by that object is not removed, as that is already assumed to have been dissipated as waste heat.
- If a nonlinear static or staged-construction load case continues from a direct-integration time-history load case, the velocity-based viscous damping energies are held constant during the static analysis. However, kinetic energy is set to zero during the static analysis and the input energy is correspondingly reduced. If a nonlinear direct-integration time-history load case then continues from the nonlinear static or staged-construction load case, the velocity continues from the end of the most recent time-history load case in the sequence and the kinetic energy is recomputed from those velocities. In effect, the nonlinear static or staged-construction load cases are assumed to have occurred instantaneously from the dynamics point of view. If this is not desired, the previous time-history load cases should be run long enough to damp out the velocities.
- The presence of energy error does not invalidate the analysis results, but can be an indication that further investigation is warranted. See Limitations below for some of the possible reasons. Note that the energy response is calculated from the basic analysis results (displacements, forces, stresses) and does not affect these results.

Limitations

Following are some limitations of the current implementation of the global energy response:

- For SAP2000 and CSiBridge, energy response is currently limited to linear and nonlinear modal time-history load cases. Nonlinear static, staged-construction, and direct-integration time-history load cases are available only for ETABS at the present time.
- Nonlinear modal time-history (FNA) load cases only support nonlinearity in link elements, hence nonlinear hysteretic damping energy from frame hinges and shear walls is only produced if the “Model Hinges as Separate Link Elements” option is used (see Topic “Analysis Modeling” in Chapter “Hinge Properties”). When using the “Model Hinges within Elements” option, hysteretic damping energy from frame hinges and shear walls is not considered. Hysteretic behavior for layered shells is not considered for FNA load cases either way.

- The input energy and potential energy due to span loads within a frame element may be under- or over-estimated if the internal deflection of the beam is significant compared to the displacements of the two ends. In such cases, subdividing or meshing the frame element may result in a different and often more accurate input and potential energy due to span loads. This effect may be especially pronounced when nonlinear hinges without hinge overwrites are present in a frame element with large span loads. Using shorter hinge overwrites or the “Model Hinges as Separate Link Elements” analysis option can be used to resolve this issue. This is a limitation on the reported energy response and does not affect any force, stress, displacement, or deformation response.
- Temperature and strain loading create self-equilibrating forces acting at the joints of the element. The input energy for this type of loading is calculated as the work done by these forces acting through the joint displacements. For a fully restrained element, no input energy would be computed, and correspondingly the potential energy will also be reported as zero. On the other hand, an unrestrained element will report non-zero work and potential energy due to temperature loading. Note that the layered shell and frame hinges behave oppositely in this regard. Their energy is computed based on the internal stress state, and thus they report nonzero input and potential energy due to restrained temperature loading, and zero energy for unrestrained loading.
- Hysteretic materials may sometimes underestimate the amount of potential energy present because they assume that unloading will occur along the original elastic stiffness, when it may actually occur along a reduced stiffness. As a result, hysteretic energy may show a slight decrease when an element unloads since some additional potential energy is recovered. This effect is generally insignificant for practical models.

Equation Solvers

You can choose from three solvers available for solving the stiffness equations:

- **Standard Solver.** This is the slowest of the solvers for larger problems and uses the most disk space. However, it provides detailed warning messages about instabilities and sensitivity of the stiffness matrix.
- **Advanced Solver.** This is the default solver. It may be one or two orders of magnitude faster than the Standard solver for larger problems, and it also uses much less disk space. It provides limited information about instabilities.
- **Multi-threaded Solver.** This is the faster solver on machines with multiple cores or processors. It runs in memory and does not use any disk space. It pro-

vides no information on instabilities. Depending on the amount of memory available, you may need to use the Advanced solver for the largest problems.

Because the solvers perform numerical operations in a different order, it is possible that sensitive problems may yield slightly different results with the different solvers due to numerical roundoff. Parallel processing can increase this effect. In extremely sensitive, nonlinear, history-dependent problems, the differences may be more pronounced.

All verification examples have been run and checked using all three solvers. The results using the advanced solver are reported for comparison in the verification manuals.

The advanced solver is based on proprietary CSI technology. It uses, in part, code derived from TAUCS family of solvers. Please see the copyright notice at the end of Chapter “References” (page 533) for more information.

Environment Variables to Control Analysis

Several environment variables are available to control the solution process. These are rarely needed. If you do use them, they can be set in the Windows Control Panel or by using the SET command in a command prompt window. In general we recommend not changing these.

Using Windows Control Panel, you can add, remove, or change the environment variables. By default they do not exist. Setting the value to zero or deleting the environment variable will result in default behavior. Changes made in the Control Panel only affect instances of SAP2000 or other CSI products that are started after the change is made. You cannot change the behavior of an instance that is already started.

Using the command prompt, changes made to an environment variable only affect instances of SAP2000 or other CSI products that are started from within that same command prompt.

SAPFIRE_NUM_THREADS

Certain parts of the analysis process are parallelized, most notably moving-load analysis and the calculation of certain element events, states, and stiffness matrices. More parallelism will be implemented over time in future versions. By default the analysis uses all the available physical cores, which may be half of the hyper-threaded cores.

You can set `SAPFIRE_NUM_THREADS` to a positive number to reduce the number of threads if you plan to run multiple instances of the analysis, or want to change the settings to obtain optimum performance from your machine.

The Multi-threaded solver uses its own logic to obtain optimum performance. However, `SAPFIRE_NUM_THREADS` does set the maximum number of threads to be used. You can override this value with the separate environment variable `SAPFIRE_NUM_THREADS_SOLVE`. This is rarely needed but could be used to achieve reproducible results in very sensitive models where the numerical calculations may be affected by the order of parallelized operations.

For the most complete control of the Multi-threaded solver, set the environment variable `MKL_DYNAMIC` to value `FALSE`, which forces the solver to use the specified number of threads. By default this value is `TRUE`. We do not generally recommend changing the default behavior, since it may affect the efficiency of the solver and other parallelized calculations. It could also affect other software products performing scientific and engineering computation if set in Windows Control Panel.

SAPFIRE_FILESIZE_MB

By default the file-size limit for internal files used during analysis is 100 GB, but this can be changed using the environment variable `SAPFIRE_FILESIZE_MB` to specify an integer number of megabytes (MB). The value can represent 1 TB or more, although the default value will rarely need to be increased. For most practical problems, the amount of disk space and other system resources will control the size of the model and associated analysis files that can be handled.

On the other hand, you will need to decrease this limit for Windows FAT32 file systems that are limited to 2 GB. In this case, set `SAPFIRE_FILESIZE_MB` to 2048 or less. Most Windows machines use NTFS file systems and are not affected by this limit. This value may also be needed for files that are intended to be zipped-up (compressed) for storage or shipping.

Accessing the Assembled Stiffness and Mass Matrices

When using the advanced equation solver, you may request that the program produce the assembled stiffness and mass matrices in the form of text files. This can be done for a single linear static, modal, or buckling Load Case. To get the stiffness and mass matrices for a nonlinear case, define a linear case that uses the stiffness from the final state of the desired nonlinear case.

The assembled matrices are provided in five text files that have the same name as the model file, but with the following extensions and contents:

- **Extension .TXA:** This file includes the counts of the number of joints and equations in the model, and also describes the format and contents of the other four files.
- **Extension .TXE:** This file gives the equation numbers for each degree of freedom (DOF) at each joint. Equation numbers are positive for active DOF that are present in the stiffness and mass matrices, negative for constrained DOF that are computed as linear combinations of active DOF, and zero for restrained or null DOF.
- **Extension .TXC:** This file defines the constraint equations, and is only present if there are constraints in the model.
- **Extension .TXK:** This file gives the lower half of the symmetric stiffness matrix.
- **Extension .TXM:** This file gives the lower half of the symmetric mass matrix.

Each of the latter four files contains a single header line that begins with “Note:” and defines the data columns. All subsequent lines provide Tab-delimited data for easy import into text editors or spreadsheet programs.

Modal Analysis

Modal analysis is used to determine the vibration modes of a structure. These modes are useful to understand the behavior of the structure. They can also be used as the basis for modal superposition in response-spectrum and modal time-history Load Cases.

Basic Topics for All Users

- Overview
- Eigenvector Analysis
- Ritz-Vector Analysis
- Modal Analysis Output

Overview

A modal analysis is defined by creating a Load Case and setting its type to “Modal”. You can define multiple modal Load Cases, resulting in multiple sets of modes.

There are two types of modal analysis to choose from when defining a modal Load Case:

- **Eigenvector** analysis determines the undamped free-vibration mode shapes and frequencies of the system. These natural modes provide an excellent insight into the behavior of the structure.
- **Ritz-vector** analysis seeks to find modes that are excited by a particular loading. Ritz vectors can provide a better basis than do eigenvectors when used for response-spectrum or time-history analyses that are based on modal superposition

You can request that static correction modes be calculated along with eigenvectors. They are automatically included with Ritz vectors. *Static correction-modes can be very important for getting accurate response at stiff supports. Their use is generally recommended.*

Modal analysis is always linear. A modal Load Case may be based on the stiffness of the full unstressed structure, or upon the stiffness at the end of a nonlinear Load Case (nonlinear static or nonlinear direct-integration time-history).

By using the stiffness at the end of a nonlinear case, you can evaluate the modes under P-delta or geometric stiffening conditions, at different stages of construction, or following a significant nonlinear excursion in a large earthquake.

In addition, the mass used for a modal analysis can be changed if multiple Mass Sources have been defined. The Mass Source to be used for a given modal load case can be specified in the same types of nonlinear Load Cases that can be used to calculate the stiffness.

For more information:

- See Topic “Mass Source” (page 334) in Chapter “Load Patterns”.
- See Chapter “Load Cases” (page 341).

Eigenvector Analysis

Eigenvector analysis determines the undamped free-vibration mode shapes and frequencies of the system. These natural Modes provide an excellent insight into the behavior of the structure. They can also be used as the basis for response-spectrum or time-history analyses, although Ritz vectors are recommended for this purpose.

Eigenvector analysis involves the solution of the generalized eigenvalue problem:

$$[K - \omega^2 M] \mathbf{0}$$

where \mathbf{K} is the stiffness matrix, \mathbf{M} is the diagonal mass matrix, ω^2 is the diagonal matrix of eigenvalues, and \mathbf{U} is the matrix of corresponding eigenvectors (mode shapes).

Each eigenvalue-eigenvector pair is called a natural Vibration Mode of the structure. The Modes are identified by numbers from 1 to n in the order in which the modes are found by the program.

The eigenvalue is the square of the circular frequency, ω , for that Mode (unless a frequency shift is used, see below). The cyclic frequency, f , and period, T , of the Mode are related to ω by:

$$T = \frac{1}{f} \quad \text{and} \quad f = \frac{\omega}{2\pi}$$

You may specify the number of modes to be found, a convergence tolerance, and the frequency range of interest. These parameters are described in the following subtopics.

Number of Modes

You may specify the maximum and minimum number of modes to be found.

The program will not calculate more than the specified maximum number of modes. This number includes any static correction modes requested. The program may compute fewer modes if there are fewer mass degrees of freedom, all dynamic participation targets have been met, or all modes within the cutoff frequency range have been found.

The program will not calculate fewer than the specified minimum number of modes, unless there are fewer mass degrees of freedom in the model.

A mass degree of freedom is any *active* degree of freedom that possesses translational mass or rotational mass moment of inertia. The mass may have been assigned directly to the joint or may come from connected elements.

Only the modes that are actually found will be available for use by any subsequent response-spectrum or modal time-history Load Cases.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom.”

Frequency Range

You may specify a restricted frequency range in which to seek the Vibration Modes by using the parameters:

- **shift**: The center of the cyclic frequency range, known as the shift frequency
- **cut**: The radius of the cyclic frequency range, known as the cutoff frequency

The program will only seek Modes with frequencies f that satisfy:

$$|f - \text{shift}| \leq \text{cut}$$

The default value of **cut** = 0 does not restrict the frequency range of the Modes.

Modes are found in order of increasing distance of frequency from the shift. This continues until the cutoff is reached, the requested number of Modes is found, or the number of mass degrees of freedom is reached.

A stable structure will possess all positive natural frequencies. When performing a seismic analysis and most other dynamic analyses, the lower-frequency modes are usually of most interest. It is then appropriate to the default shift of zero, resulting in the lowest-frequency modes of the structure being calculated. If the shift is not zero, response-spectrum and time-history analyses may be performed; however, static, moving-load, and p-delta analyses are not allowed.

If the dynamic loading is known to be of high frequency, such as that caused by vibrating machinery, it may be most efficient to use a positive shift near the center of the frequency range of the loading.

A structure that is unstable when unloaded will have some modes with zero frequency. These modes may correspond to rigid-body motion of an inadequately supported structure, or to mechanisms that may be present within the structure. It is not possible to compute the static response of such a structure. However, by using a small negative shift, the lowest-frequency vibration modes of the structure, including the zero-frequency instability modes, can be found. This does require some mass to be present that is activated by each instability mode.

A structure that has buckled under P-delta load will have some modes with zero or negative frequency. During equation solution, the number of frequencies less than the shift is determined and printed in the log file. If you are using a zero or negative shift and the program detects a negative-frequency mode, it will stop the analysis since the results will be meaningless. If you use a positive shift, the program will permit negative frequencies to be found; however, subsequent static and dynamic results are still meaningless.

When using a frequency shift, the stiffness matrix is modified by subtracting from it the mass matrix multiplied by ω_0^2 , where $\omega_0 = 2 \pi \text{ shift}$. If the shift is very near a natural frequency of the structure, the solution becomes unstable and will complain during equation solution. Run the analysis again using a slightly different shift frequency.

The circular frequency, ω , of a Vibration Mode is determined from the shifted eigenvalue, λ , as:

$$\omega = \sqrt{\lambda + \omega_0^2}$$

Automatic Shifting

As an option, you may request that the eigen-solver use automatic shifting to speed up the solution and improve the accuracy of the results. This is particularly helpful when seeking a large number of modes, for very large structures, or when there are a lot of closely spaced modes to be found.

The solver will start with the requested shift frequency, **shift** (default zero), and then successively then shift to the right (in the positive direction) as needed to improve the rate of convergence.

If no cutoff frequency has been specified (**cut** = 0), automatic shifting will only be to the right, which means that eigenvalues to the left of the initial shift may be missed. This is not usually a problem for stable structures starting with an initial shift of zero.

If a cutoff frequency has been specified (**cut** > 0), automatic shifting will be to the right until all eigenvalues between **shift** and **shift** + **cut** have been found, then the automatic shifting will return to the initial shift and proceed to the left from there.

In either case, automatic shifting may not find eigenvalues in the usual order of increasing distance from the initial shift.

Convergence Tolerance

SAP2000 solves for the eigenvalue-eigenvectors pairs using an accelerated subspace iteration algorithm. During the solution phase, the program prints the approximate eigenvalues after each iteration. As the eigenvectors converge they are removed from the subspace and new approximate vectors are introduced. For details of the algorithm, see Wilson and Tetsuji (1983).

You may specify the relative convergence tolerance, **tol**, to control the solution; the default value is $\text{tol} = 10^{-9}$. The iteration for a particular Mode will continue until the relative change in the eigenvalue between successive iterations is less than 2 tol , that is until:

$$\frac{1}{2} \frac{\lambda_{i+1} - \lambda_i}{\lambda_{i+1}} \leq \text{tol}$$

where λ is the eigenvalue relative to the frequency shift, and i and $i + 1$ denote successive iteration numbers.

In the usual case where the frequency shift is zero, the test for convergence becomes approximately the same as:

$$\frac{T_{i+1} - T_i}{T_{i+1}} \leq \text{tol} \quad \text{or} \quad \frac{f_{i+1} - f_i}{f_i} \leq \text{tol}$$

provided that the difference between the two iterations is small.

Note that the error in the eigenvectors will generally be larger than the error in the eigenvalues. The relative error in the global force balance for a given Mode gives a measure of the error in the eigenvector. This error can usually be reduced by using a smaller value of **tol**, at the expense of more computation time.

Static-Correction Modes

Static correction-modes can be very important for getting accurate response at stiff supports. Their use is generally recommended.

You may request that the program compute the static-correction modes for any Acceleration Load or Load Pattern. A static-correction mode is the static solution to that portion of the specified load that is not represented by the found eigenvectors.

When applied to acceleration loads, static-correction modes are also known as missing-mass modes or residual-mass modes.

Static-correction modes are of little interest in their own right. They are intended to be used as part of a modal basis for response-spectrum or modal time-history analysis for high frequency loading to which the structure responds statically. Although a static-correction mode will have a mode shape and frequency (period) like the eigenvectors do, it is not a true eigenvector.

You can specify for which Load Patterns and/or Acceleration Loads you want static correction modes calculated, if any. One static-correction mode will be computed for each specified Load unless all eigenvectors that can be excited by that Load have been found. Static-correction modes count against the maximum number of modes requested for the Load Case.

As an example, consider the translational acceleration load in the UX direction, m_x . Define the participation factor for mode n as:

$$f_{xn} = \mathbf{m}_x^T \mathbf{m}_n$$

The static-correction load for UX translational acceleration is then:

$$\mathbf{m}_{x0} = \mathbf{m}_x \sum_{n=1}^N \frac{f_{xn}}{M_n}$$

The static-correction mode-shape vector, \mathbf{m}_{x0} , is the solution to:

$$\mathbf{K} \mathbf{m}_{x0} = \mathbf{m}_{x0}$$

If \mathbf{m}_{x0} is found to be zero, all of the modes necessary to represent UX acceleration have been found, and no residual-mass mode is needed or will be calculated.

The static-correction modes for any other acceleration load or Load Pattern are defined similarly.

Each static-correction mode is assigned a frequency that is calculated using the standard Rayleigh quotient method. When static-correction modes are calculated, they are used for Response-spectrum and Time-history analysis just as the eigenvectors are.

The use of static-correction modes assures that the static-load participation ratio will be 100% for the selected acceleration loads. *However, static-correction modes do not generally result in mass-participation ratios or dynamic-load participation ratios of 100%.* Only true dynamic modes (eigen or Ritz vectors) can increase these ratios to 100%.

See Topic “Modal Analysis Output” (page 349) in this Chapter for more information on modal participation ratios.

Note that Ritz vectors, described next, always include the residual-mass effect for all starting load vectors.

Ritz-Vector Analysis

Research has indicated that the natural free-vibration mode shapes are not the best basis for a mode-superposition analysis of structures subjected to dynamic loads. It has been demonstrated (Wilson, Yuan, and Dickens, 1982) that dynamic analyses based on a special set of load-dependent Ritz vectors yield more accurate results than the use of the same number of natural mode shapes. The algorithm is detailed in Wilson (1985).

The reason the Ritz vectors yield excellent results is that they are generated by taking into account the spatial distribution of the dynamic loading, whereas the direct use of the natural mode shapes neglects this very important information.

In addition, the Ritz-vector algorithm automatically includes the advantages of the proven numerical techniques of static condensation, Guyan reduction, and static correction due to higher-mode truncation.

The spatial distribution of the dynamic load vector serves as a **starting load vector** to initiate the procedure. The first Ritz vector is the static displacement vector corresponding to the starting load vector. The remaining vectors are generated from a recurrence relationship in which the mass matrix is multiplied by the previously obtained Ritz vector and used as the load vector for the next static solution. Each static solution is called a **generation cycle**.

When the dynamic load is made up of several independent spatial distributions, each of these may serve as a starting load vector to generate a set of Ritz vectors. Each generation cycle creates as many Ritz vectors as there are starting load vectors. If a generated Ritz vector is redundant or does not excite any mass degrees of freedom, it is discarded and the corresponding starting load vector is removed from all subsequent generation cycles.

Standard eigen-solution techniques are used to orthogonalize the set of generated Ritz vectors, resulting in a final set of Ritz-vector Modes. Each Ritz-vector Mode consists of a mode shape and frequency. The full set of Ritz-vector Modes can be used as a basis to represent the dynamic displacement of the structure.

When a sufficient number of Ritz-vector Modes have been found, some of them may closely approximate natural mode shapes and frequencies. In general, however, Ritz-vector Modes do not represent the intrinsic characteristics of the structure in the same way the natural Modes do. The Ritz-vector modes are biased by the starting load vectors.

You may specify the number of Modes to be found, the starting load vectors to be used, and the number of generation cycles to be performed for each starting load vector. These parameters are described in the following subtopics.

Number of Modes

You may specify the maximum and minimum number of modes to be found.

The program will not calculate more than the specified maximum number of modes. The program may compute fewer modes if there are fewer mass degrees of freedom, all dynamic participation targets have been met, or the maximum number of cycles has been reached for all loads.

The program will not calculate fewer than the specified minimum number of modes, unless there are fewer mass degrees of freedom in the model.

A mass degree of freedom is any *active* degree of freedom that possesses translational mass or rotational mass moment of inertia. The mass may have been assigned directly to the joint or may come from connected elements.

Only the modes that are actually found will be available for use by any subsequent response-spectrum or modal time-history Load Cases.

See Topic “Degrees of Freedom” (page 30) in Chapter “Joints and Degrees of Freedom.”

Starting Load Vectors

You may specify any number of starting load vectors. Each starting load vector may be one of the following:

- An Acceleration Load in the global X, Y, or Z direction
- A Load Pattern
- A built-in nonlinear deformation load, as described below

For response-spectrum analysis, only the Acceleration Loads are needed. For modal time-history analysis, one starting load vector is needed for each Load Pattern or Acceleration Load that is used in any modal time-history.

If nonlinear modal time-history analysis is to be performed, an additional starting load vector is needed for *each* independent nonlinear deformation. You may specify that the program use the built-in nonlinear deformation loads, or you may define your own Load Patterns for this purpose. See Topic “Nonlinear Deformation

Loads” (page 271) in Chapter “The Link/Support Element—Basic” for more information.

If you define your own starting load vectors, do the following for *each* nonlinear deformation:

- Explicitly define a Load Pattern that consists of a set of self-equilibrating forces that activates the desired nonlinear deformation
- Specify that Load Pattern as a starting load vector

The number of such Load Patterns required is equal to the number of independent nonlinear deformations in the model.

If several Link/Support elements act together, you may be able to use fewer starting load vectors. For example, suppose the horizontal motion of several base isolators are coupled with a diaphragm. Only three starting load vectors acting on the diaphragm are required: two perpendicular horizontal loads and one moment about the vertical axis. Independent Load Cases may still be required to represent any vertical motions or rotations about the horizontal axes for these isolators.

It is *strongly recommended* that mass (or mass moment of inertia) be present at every degree of freedom that is loaded by a starting load vector. This is automatic for Acceleration Loads, since the load is caused by mass. If a Load Pattern or nonlinear deformation load acts on a non-mass degree of freedom, the program issues a warning. Such starting load vectors may generate inaccurate Ritz vectors, or even no Ritz vectors at all.

Generally, the more starting load vectors used, the more Ritz vectors must be requested to cover the same frequency range. Thus including unnecessary starting load vectors is not recommended.

In each generation cycle, Ritz vectors are found in the order in which the starting load vectors are specified. In the last generation cycle, only as many Ritz vectors will be found as required to reach the total number of Modes, **n**. For this reason, the most important starting load vectors should be specified first, especially if the number of starting load vectors is not much smaller than the total number of Modes.

For more information:

- See Topic “Nonlinear Modal Time-History Analysis (FNA)” (page 149) in Chapter “Nonlinear Time-History Analysis”.
- See Chapter “Load Patterns” (page 321).

Number of Generation Cycles

You may specify the maximum number of generation cycles, **ncyc**, to be performed for each starting load vector. This enables you to obtain more Ritz vectors for some starting load vectors than others. By default, the number of generation cycles performed for each starting load vector is unlimited, i.e., until the total number, **n**, of requested Ritz vectors have been found.

As an example, suppose that two linear time-history analyses are to be performed:

- (1) Gravity load is applied quasi-statically to the structure using Load Patterns DL and LL
- (2) Seismic load is applied in all three global directions

The starting load vectors required are the three Acceleration Loads and Load Patterns DL and LL. The first generation cycle creates the static solution for each starting load vector. This is all that is required for Load Patterns DL and LL in the first History, hence for these starting load vectors **ncyc** = 1 should be specified. Additional Modes may be required to represent the dynamic response to the seismic loading, hence an unlimited number of cycles should be specified for these starting load vectors. If 12 Modes are requested (**n** = 12), there will be one each for DL and LL, three each for two of the Acceleration Loads, and four for the Acceleration Load that was specified first as a starting load vector.

Starting load vectors corresponding to nonlinear deformation loads may often need only a limited number of generation cycles. Many of these loads affect only a small local region and excite only high-frequency natural modes that may respond quasi-statically to typical seismic excitation. If this is the case, you may be able to specify **ncyc** = 1 or 2 for these starting load vectors. More cycles may be required if you are particularly interested in the dynamic behavior in the local region.

You must use your own engineering judgment to determine the number of Ritz vectors to be generated for each starting load vector. No simple rule can apply to all cases.

Modal Analysis Output

Various properties of the Vibration Modes are available as analysis results. This information is the same regardless of whether you use eigenvector or Ritz-vector analysis, and is described in the following subtopics.

Periods and Frequencies

The following time-properties are printed for each Mode:

- Period, T , in units of time
- Cyclic frequency, f , in units of cycles per time; this is the inverse of T
- Circular frequency, ω , in units of radians per time; $\omega = 2\pi f$
- Eigenvalue, ω^2 , in units of radians-per-time squared

Participation Factors

The modal participation factors are the dot products of the six Acceleration Loads with the modes shapes. The participation factors for Mode n corresponding to translational Acceleration Loads in the global X, Y, and Z directions are given by:

$$f_{xn} = \mathbf{m}_x^T \mathbf{m}_n$$

$$f_{yn} = \mathbf{m}_y^T \mathbf{m}_n$$

$$f_{zn} = \mathbf{m}_z^T \mathbf{m}_n$$

where \mathbf{m}_n is the mode shape and \mathbf{m}_x , \mathbf{m}_y , and \mathbf{m}_z are the unit translational Acceleration Loads.

Similarly, the participation factors corresponding to rotational Acceleration Loads about the centroidal axes parallel to the global X, Y, and Z axes are given by:

$$f_{rxn} = \mathbf{m}_{rx}^T \mathbf{m}_n$$

$$f_{ryn} = \mathbf{m}_{ry}^T \mathbf{m}_n$$

$$f_{rzn} = \mathbf{m}_{rz}^T \mathbf{m}_n$$

Here \mathbf{m}_{rx} , \mathbf{m}_{ry} , and \mathbf{m}_{rz} are the unit rotational Acceleration Loads.

These factors are the generalized loads acting on the Mode due to each of the Acceleration Loads. These values are called “factors” because they are related to the mode shape and to a unit acceleration. The modes shapes are each normalized, or scaled, with respect to the mass matrix such that:

$$\mathbf{m}_n^T \mathbf{M} \mathbf{m}_n = 1$$

The actual magnitudes and signs of the participation factors are not important. What is important is the relative values of the six factors for a given Mode.

Important: Although the rotational accelerations are applied in load cases about the origins of the specified coordinate systems, the modal participation factors and the various modal participation ratios described below for the rotational accelerations are reported about the center of mass for the structure. This makes the rotational participation factors and ratios more meaningful because they do not include any contribution from the translational accelerations

For more information, See Topic “Acceleration Loads” (page 338) in Chapter “Load Patterns.”

Participating Mass Ratios

The participating mass ratio for a Mode provides a measure of how important the Mode is for computing the response to the six Acceleration Loads in the global coordinate system. Thus it is useful for determining the accuracy of response-spectrum analyses and seismic time-history analyses. The participating mass ratio provides *no information* about the accuracy of time-history analyses subjected to other loads.

The participating mass ratios for Mode n corresponding to translational Acceleration Loads in the global X, Y, and Z directions are given by:

$$r_{xn} = \frac{(f_{xn})^2}{M_x}$$

$$r_{yn} = \frac{(f_{yn})^2}{M_y}$$

$$r_{zn} = \frac{(f_{zn})^2}{M_z}$$

where f_{xn} , f_{yn} , and f_{zn} are the participation factors defined in the previous subtopic; and M_x , M_y , and M_z are the total unrestrained masses acting in the global X, Y, and Z directions.

The participating mass ratios corresponding to rotational Acceleration Loads about centroidal axes parallel to the global X, Y, and Z directions are given by:

$$r_{rxn} = \frac{(f_{rxn})^2}{M_{rx}}$$

$$r_{ryy} = \frac{(f_{ryy})^2}{M_{ry}}$$

$$r_{rzn} = \frac{(f_{rzn})^2}{M_{rz}}$$

where f_{rxn} , f_{ryy} , and f_{rzn} are the participation factors defined in the previous subtopic; and M_{rx} , M_{ry} , and M_{rz} are the total rotational inertias of the unrestrained masses acting about the centroidal axes parallel to the global X, Y, and Z directions.

The cumulative sums of the participating mass ratios for all Modes up to Mode n are printed with the individual values for Mode n . This provides a simple measure of how many Modes are required to achieve a given level of accuracy for ground-acceleration loading.

If all eigen Modes of the structure are present, the participating mass ratio for each of the Acceleration Loads should generally be unity (100%). However, this may not be the case in the presence of Asolid elements or certain types of Constraints where symmetry conditions prevent some of the mass from responding to translational accelerations.

Static and Dynamic Load Participation Ratios

The static and dynamic load participation ratios provide a measure of how adequate the calculated modes are for representing the response to time-history analyses. These two measures are printed in the output file for each of the following spatial load vectors:

- The three unit Acceleration Loads
- Three rotational Acceleration loads
- All Load Patterns specified in the definition of the modal Load Case
- All nonlinear deformation loads, if they are specified in the definition of the modal Load Case

The Load Patterns and Acceleration Loads represent spatial loads that you can explicitly specify in a modal time-history analysis, whereas the last represents loads that can act implicitly in a *nonlinear* modal time-history analysis.

For more information:

- See Topic “Nonlinear Deformation Loads” (page 271) in Chapter “The Link/Support Element—Basic.”
- See Chapter “Load Patterns” (page 321).
- See Topic “Acceleration Loads” (page 338) in Chapter “Load Patterns.”
- See Topic “Linear Modal Time-History Analysis” (page 403) in Chapter “Linear Time-History Analysis”.
- See Topic “Nonlinear Modal Time-History Analysis” (page 149) in Chapter “Nonlinear Time-History Analysis”.

Static Load Participation Ratio

The static load participation ratio measures how well the calculated modes can represent the response to a given static load. This measure was first presented by Wilson (1997). For a given spatial load vector \mathbf{p} , the participation factor for Mode n is given by

$$f_n = \mathbf{u}_n^T \mathbf{p}$$

where \mathbf{u}_n is the mode shape (vector) of Mode n . This factor is the generalized load acting on the Mode due to load \mathbf{p} . Note that f_n is just the usual participation factor when \mathbf{p} is one of the six unit Acceleration Loads.

The static participation ratio for this mode is given by

$$r_n^S = \frac{f_n^2}{\mathbf{u}^T \mathbf{p}}$$

where \mathbf{u} is the static solution given by $\mathbf{K}\mathbf{u} = \mathbf{p}$. This ratio gives the fraction of the total strain energy in the exact static solution that is contained in Mode n . Note that the denominator can also be represented as $\mathbf{u}^T \mathbf{K}\mathbf{u}$.

Finally, the cumulative sum of the static participation ratios for all the calculated modes is printed in the output file:

$$R^S = \sum_{n=1}^N r_n^S = \frac{\sum_{n=1}^N \mathbf{u}_n^T \mathbf{p}}{\mathbf{u}^T \mathbf{p}}$$

where N is the number of modes found. This value gives the fraction of the total strain energy in the exact static solution that is captured by the N modes.

When solving for static solutions using quasi-static time-history analysis, the value of R^S should be close to 100% for any applied static Loads, and also for all nonlinear deformation loads if the analysis is nonlinear.

Note that when Ritz-vectors are used, the value of R^S will always be 100% for all starting load vectors. This may not be true when eigenvectors are used without static correction modes. In fact, even using all possible eigenvectors will not give 100% static participation if load \mathbf{p} acts on any massless degrees-of-freedom, or if the system is sensitive or ill-conditioned. Static-correction or Ritz modes are highly recommended in these cases.

Dynamic Load Participation Ratio

The dynamic load participation ratio measures how well the calculated modes can represent the response to a given dynamic load. This measure was developed for SAP2000, and it is an extension of the concept of participating mass ratios. It is assumed that the load acts only on degrees of freedom with mass. Any portion of load vector \mathbf{p} that acts on massless degrees of freedom cannot be represented by this measure and is ignored in the following discussion.

For a given spatial load vector \mathbf{p} , the participation factor for Mode n is given by

$$f_n = \mathbf{a}_n^T \mathbf{p}$$

where \mathbf{a}_n is the mode shape for Mode n . Note that f_n is just the usual participation factor when \mathbf{p} is one of the six unit Acceleration Loads.

The dynamic participation ratio for this mode is given by

$$r_n^D = \frac{f_n^2}{\mathbf{a}^T \mathbf{p}}$$

where \mathbf{a} is the acceleration given by $\mathbf{M}\mathbf{a} = \mathbf{p}$. The acceleration \mathbf{a} is easy to calculate since \mathbf{M} is diagonal. The values of \mathbf{a} and \mathbf{p} are taken to be zero at all massless degrees of freedom. Note that the denominator can also be represented as $\mathbf{a}^T \mathbf{M}\mathbf{a}$.

Finally, the cumulative sum of the dynamic participation ratios for all the calculated modes is printed in the output file:

$$R^D = \sum_{n=1}^N r_n^D = \frac{\sum_{n=1}^N \mathbf{a}_n^T \mathbf{p}}{\mathbf{a}^T \mathbf{p}}$$

where N is the number of modes found. When \mathbf{p} is one of the unit acceleration loads, r^D is the usual mass participation ratio, and R^D is the usual cumulative mass participation ratio.

When R^D is 100%, the calculated modes should be capable of exactly representing the solution to any time-varying application of spatial load \mathbf{p} . If R^D is less than 100%, the accuracy of the solution will depend upon the frequency content of the time-function multiplying load \mathbf{p} . Normally it is the high frequency response that is not captured when R^D is less than 100%.

The dynamic load participation ratio only measures how the modes capture the spatial characteristics of \mathbf{p} , not its temporal characteristics. For this reason, R^D serves only as a qualitative guide as to whether enough modes have been computed. You must still examine the response to each different dynamic loading with varying number of modes to see if enough modes have been used.

Response-Spectrum Analysis

Response-spectrum analysis is a statistical type of analysis for the determination of the likely response of a structure to seismic loading.

Basic Topics for All Users

- Overview
- Local Coordinate System
- Response-Spectrum Function
- Modal Damping
- Modal Combination
- Directional Combination
- Response-Spectrum Analysis Output

Overview

The dynamic equilibrium equations associated with the response of a structure to ground motion are given by:

$$K \mathbf{u}(t) + C \dot{\mathbf{u}}(t) + M \ddot{\mathbf{u}}(t) = \mathbf{m}_x \ddot{u}_{gx}(t) + \mathbf{m}_y \ddot{u}_{gy}(t) + \mathbf{m}_z \ddot{u}_{gz}(t)$$

where \mathbf{K} is the stiffness matrix; \mathbf{C} is the proportional damping matrix; \mathbf{M} is the diagonal mass matrix; \mathbf{u} , $\dot{\mathbf{u}}$, and $\ddot{\mathbf{u}}$ are the relative displacements, velocities, and accelerations with respect to the ground; \mathbf{m}_x , \mathbf{m}_y , and \mathbf{m}_z are the unit Acceleration Loads; and \ddot{u}_{gx} , \ddot{u}_{gy} , and \ddot{u}_{gz} are the components of uniform ground acceleration.

Response-spectrum analysis seeks the likely maximum response to these equations rather than the full time history. The earthquake ground acceleration in each direction is given as a digitized response-spectrum curve of pseudo-spectral acceleration response versus period of the structure.

Even though accelerations may be specified in three directions, only a single, positive result is produced for each response quantity. The response quantities include displacements, forces, and stresses. Each computed result represents a statistical measure of the likely maximum magnitude for that response quantity. The actual response can be expected to vary within a range from this positive value to its negative.

No correspondence between two different response quantities is available. No information is available as to when this extreme value occurs during the seismic loading, or as to what the values of other response quantities are at that time.

Response-spectrum analysis is performed using mode superposition (Wilson and Button, 1982). Modes may have been computed using eigenvector analysis or Ritz-vector analysis. Ritz vectors are recommended since they give more accurate results for the same number of Modes. You must define a Modal Load Case that computes the modes, and then refer to that Modal Load Case in the definition of the Response-Spectrum Case.

Response-spectrum can consider high-frequency rigid response if requested and if appropriate modes have been computed. When eigen modes are used, you should request that static correction vectors be computed. This information is automatically available in Ritz modes generated for ground acceleration. In either case, you must be sure to have sufficient dynamical modes below the rigid frequency of the ground motion.

Any number of response-spectrum Load Cases can be defined. Each case can differ in the acceleration spectra applied and in the way that results are combined. Different cases can also be based upon different sets of modes computed in different Modal Load Cases. For example, this would enable you to consider the response at different stages of construction, or to compare the results using eigenvectors and Ritz vectors.

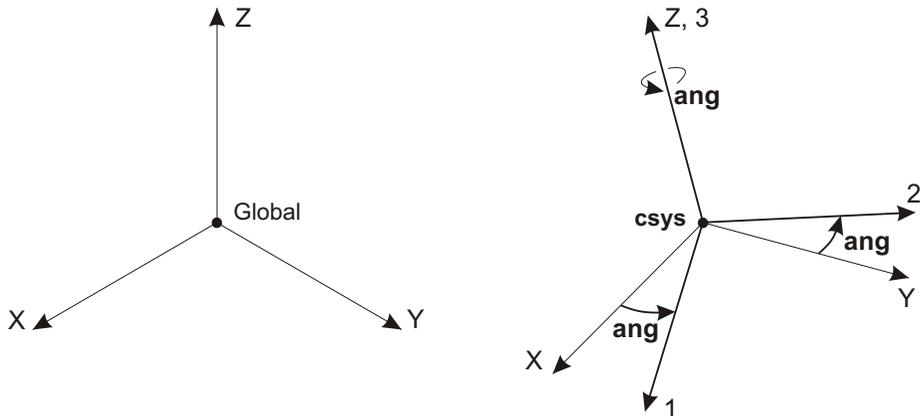


Figure 80
Definition of Response Spectrum Local Coordinate System

Local Coordinate System

Each Spec has its own **response-spectrum local coordinate system** used to define the directions of ground acceleration loading. The axes of this local system are denoted 1, 2, and 3. By default these correspond to the global X, Y, and Z directions, respectively.

You may change the orientation of the local coordinate system by specifying:

- A fixed coordinate system **csys** (the default is zero, indicating the global coordinate system)
- A coordinate angle, **ang** (the default is zero)

The local 3 axis is always the same as the Z axis of coordinate system **csys**. The local 1 and 2 axes coincide with the X and Y axes of **csys** if angle **ang** is zero. Otherwise, **ang** is the angle from the X axis to the local 1 axis, measured counterclockwise when the +Z axis is pointing toward you. This is illustrated in Figure 80 (page 385).

Response-Spectrum Function

The response-spectrum curve for a given direction is defined by digitized points of pseudo-spectral acceleration response versus period of the structure. The shape of

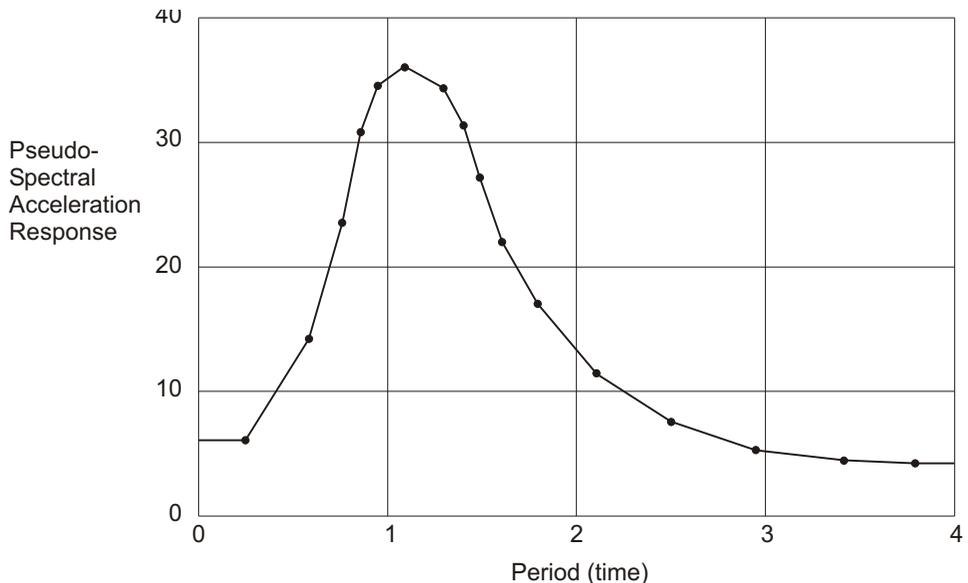


Figure 81
Digitized Response-Spectrum Curve

the curve is given by specifying the name of a Function. All values for the abscissa and ordinate of this Function must be zero or positive. See (page 381).

The function is assumed to be normalized with respect to gravity. You may specify a scale factor **sf** to multiply the ordinate (pseudo-spectral acceleration response) of the function. This should be used to convert the normalized acceleration to units consistent with the rest of the model. The scale factor itself has acceleration units and will be automatically converted if you change length units.

If the response-spectrum curve is not defined over a period range large enough to cover the Vibration Modes of the structure, the curve is extended to larger and smaller periods using a constant acceleration equal to the value at the nearest defined period.

See Topic “Functions” (page 350) in this Chapter for more information.

Damping

The response-spectrum curve chosen should reflect the damping that is present in the structure being modeled. Note that the damping is inherent in the shape of the

response-spectrum curve itself. As part of the response-spectrum function definition, you must specify the damping value that was used to generate the response-spectrum curve. During the analysis, the response-spectrum curve will automatically be adjusted from this damping value to the actual damping present in the model.

If zero damping is specified for either the response-spectrum function or the response-spectrum load case, no scaling will be performed.

Modal Damping

Damping in the structure has two effects on response-spectrum analysis:

- It modifies the shape of the response-spectrum input curve
- It affects the amount of statistical coupling between the modes for certain methods of response-spectrum modal combination (e.g., CQC and GMC)

The damping in the structure is modeled using uncoupled modal damping. Each mode has a damping ratio, **damp**, which is measured as a fraction of critical damping and must satisfy:

$$0 \leq \text{damp} \leq 1$$

Modal damping has three different sources, which are described in the following. Damping from these sources are added together. The program automatically makes sure that the total is less than one.

Modal Damping from the Load Case

For each response-spectrum Load Case, you may specify modal damping ratios that are:

- Constant for all modes
- Linearly interpolated by period or frequency. You specify the damping ratio at a series of frequency or period points. Between specified points the damping is linearly interpolated. Outside the specified range, the damping ratio is constant at the value given for the closest specified point.
- Mass and stiffness proportional. This mimics the proportional damping used for direct-integration, except that the damping value is never allowed to exceed unity.

In addition, you may optionally specify damping overwrites. These are specific values of damping to be used for specific modes that replace the damping obtained by one of the methods above. The use of damping overwrites is rarely necessary.

Composite Modal Damping from the Materials

Modal damping ratios, if any, that have been specified for the Materials are converted automatically to composite modal damping. Any cross coupling between the modes is ignored. These modal-damping values will generally be different for each mode, depending upon how much deformation each mode causes in the elements composed of the different Materials.

Effective Damping from the Link/Support Elements

Linear effective-damping coefficients, if any, that have been specified for Link/Support elements in the model are automatically converted to modal damping. Any cross coupling between the modes is ignored. These effective modal-damping values will generally be different for each mode, depending upon how much deformation each mode causes in the Link/Support elements.

Modal Combination

For a given direction of acceleration, the maximum displacements, forces, and stresses are computed throughout the structure for each of the Vibration Modes. These modal values for a given response quantity are combined to produce a single, positive result for the given direction of acceleration. The response has two parts: periodic and rigid. You can control the contribution of these two parts by specifying controlling frequencies that are properties of the seismic loading.

In addition, you can choose the statistical method used to compute the periodic response. Modal damping, as described in the previous topic, may affect the coupling between the modes, depending upon the method chosen for periodic modal combination.

Periodic and Rigid Response

For all modal combination methods except Absolute Sum, there are two parts to the modal response for a given direction of loading: periodic and rigid. The distinction here is a property of the loading, not of the structure. Two frequencies are defined, **f1** and **f2**, which define the rigid-response content of the ground motion, where **f1** and **f2**.

For structural modes with frequencies less than **f1** (longer periods), the response is fully periodic. For structural modes with frequencies above **f2** (shorter periods), the response is fully rigid. Between frequencies **f1** and **f2**, the amount of periodic and rigid response is interpolated, as described by Gupta (1990).

Frequencies **f1** and **f2** are properties of the seismic input, not of the structure. Gupta defines **f1** as:

$$\mathbf{f1} = \frac{S_{A_{\max}}}{2 S_{V_{\max}}}$$

where $S_{A_{\max}}$ is the maximum spectral acceleration and $S_{V_{\max}}$ is the maximum spectral velocity for the ground motion considered. The default value for **f1** is unity.

Gupta defines **f2** as:

$$\mathbf{f2} = \frac{1}{3} \mathbf{f1} + \frac{2}{3} f_r$$

where f_r is the rigid frequency of the seismic input, i.e., that frequency above which the spectral acceleration is essentially constant and equal to the value at zero period (infinite frequency). Others have defined **f2** as:

$$\mathbf{f2} = f_r$$

The following rules apply when specifying **f1** and **f2**:

- If **f2** = 0, no rigid response is calculated and all response is periodic, regardless of the value specified for **f1**.
- Otherwise, the following condition must be satisfied: $0 \leq \mathbf{f1} \leq \mathbf{f2}$.
- Specifying **f1** = 0 is the same as specifying **f1** = **f2**.

For any given response quantity (displacement, stress, force, etc.), the periodic response, R_p , is computed by one of the modal combination methods described below. The rigid response, R_r , is always computed as an algebraic (fully correlated) sum of the response from each mode having frequency above **f2**, and an interpolated portion of the response from each mode between **f1** and **f2**. The total response, R , is computed by one of the following two methods:

- SRSS, as recommended by Gupta (1990) and NRC (2006), which assumes that these two parts are statistically independent:

$$R = \sqrt{R_p^2 + R_r^2}$$

- Absolute Sum, for compatibility with older methods:

$$R = |R_p| + |R_r|$$

Please note that the choice of using the SRSS or Absolute Sum for combining periodic and rigid response is independent of the periodic modal combination or the directional combination methods described below.

CQC Method

The Complete Quadratic Combination technique for calculating the periodic response is described by Wilson, Der Kiureghian, and Bayo (1981). This is the default method of modal combination.

The CQC method takes into account the statistical coupling between closely-spaced Modes caused by modal damping. Increasing the modal damping increases the coupling between closely-spaced modes. If the damping is zero for all Modes, this method degenerates to the SRSS method.

GMC Method

The General Modal Combination technique for calculating the periodic response is the complete modal combination procedure described by Equation 3.31 in Gupta (1990). The GMC method takes into account the statistical coupling between closely-spaced Modes similarly to the CQC method, but uses the Rosenblueth correlation coefficient with the time duration of the strong earthquake motion set to infinity. The result is essentially identical to the CQC method.

Increasing the modal damping increases the coupling between closely-spaced modes. If the damping is zero for all Modes, this method degenerates to the SRSS method.

SRSS Method

This method for calculating the periodic response combines the modal results by taking the square root of the sum of their squares. This method does not take into account any coupling of the modes, but rather assumes that the response of the modes are all statistically independent. Modal damping does not affect the results.

Absolute Sum Method

This method combines the modal results by taking the sum of their absolute values. Essentially all modes are assumed to be fully correlated. This method is usually over-conservative. The distinction between periodic and rigid response is not considered for this method. All modes are treated equally. Modal damping does not affect the results.

NRC Ten-Percent Method

This technique for calculating the periodic response is the Ten-Percent method of the U.S. Nuclear Regulatory Commission *Regulatory Guide 1.92* (NRC, 2006).

The Ten-Percent method assumes full, positive coupling between all modes whose frequencies differ from each other by 10% or less of the smaller of the two frequencies. Modal damping does not affect the coupling.

NRC Double-Sum Method

This technique for calculating the periodic response is the Double-Sum method of the U.S. Nuclear Regulatory Commission *Regulatory Guide 1.92*. (NRC, 2006).

The Double-Sum method assumes a positive coupling between all modes, with correlation coefficients that depend upon damping in a fashion similar to the CQC and GMC methods, and that also depend upon the duration of the earthquake. You specify this duration as parameter **td** as part of the Load Cases definition.

Directional Combination

For each displacement, force, or stress quantity in the structure, the modal combination produces a single, positive result for each direction of acceleration. These directional values for a given response quantity are combined to produce a single, positive result. Three methods are available for combining the directional response, SRSS, CQC3, and Absolute Sum.

SRSS Method

This method combines the response for different directions of loading by taking the square root of the sum of their squares:

$$R \sqrt{R_1^2 + R_2^2 + R_3^2}$$

where R_1 , R_2 , and R_3 are the modal-combination values for each direction. This method is invariant with respect to coordinate system, i.e., the results do not depend upon your choice of coordinate system when the given response-spectrum curves are the same in each direction. This is the default method for directional combination, and is closely related to the CQC3 method described next.

CQC3 Method

The CQC3 method (Menun and Der Kiureghian, 1998) is an extension of the SRSS method of directional combination. It is applicable when the two horizontal spectra are identical in shape but have different scale factors, as is often assumed. When the direction of loading for the two spectra is not known, it is necessary to consider the envelope of loading at all possible angles.

The CQC3 method does this automatically by calculating the critical loading angle for each response quantity, and reporting the maximum response at that angle. All that is required is to specify the same response-spectrum function for directions U1 and U2, but with two different scale factors, and to select the CQC3 method for directional combination. The same response will be obtained no matter what value you specify for the loading angle, **ang**, in a given coordinate system, **csys**, since all angles are enveloped.

The response to vertical loading in direction U3, if present, is combined with the maximum horizontal response using the SRSS rule. No variation of the vertical direction is considered.

If the horizontal spectra and their scale factors are both identical, the CQC3 method degenerates to the SRSS method.

If different spectra are specified for the two horizontal directions, the CQC3 method may still be selected and the same calculations will be performed. However, the results are no longer completely independent of loading angle, and they must be reviewed by an engineer for their significance.

The CQC3 method was originally defined for periodic response and for the CQC method of modal combination. It has been extended in SAP2000 to apply to all types of modal combination, and also to include the rigid response, if any. When the absolute modal combination is used, the CQC3 results are not completely independent of loading angle, but for all quadratic types of modal combination, angular independence is obtained. CQC3 can be recommended over the SRSS method un-

less the direction of loading is known. Both methods are independent of the choice of global coordinate system.

Absolute Sum Method

This method combines the response for different directions of loading by taking the sum of their absolute values. A scale factor, **dirf**, is available for reducing the interaction between the different directions.

Specify **dirf** = 1 for a simple absolute sum:

$$R = |R_1| + |R_2| + |R_3|$$

This method is usually over-conservative.

Specify $0 < \mathbf{dirf} < 1$ to combine the directional results by the scaled absolute sum method. Here, the directional results are combined by taking the maximum, over all directions, of the sum of the absolute values of the response in one direction plus **dirf** times the response in the other directions.

For example, if **dirf** = 0.3, the spectral response, R , for a given displacement, force, or stress would be:

$$R = \max(\bar{R}_1, \bar{R}_2, \bar{R}_3)$$

where:

$$\bar{R}_1 = R_1 + 0.3(R_2 + R_3)$$

$$\bar{R}_2 = R_2 + 0.3(R_1 + R_3)$$

$$\bar{R}_3 = R_3 + 0.3(R_1 + R_2)$$

and R_1 , R_2 , and R_3 are the modal-combination values for each direction.

Unlike the SRSS and CQC3 methods, the absolute sum method can give different results depending upon your arbitrary choice of coordinate system, even when the angle between the direction of loading and the principal axes of the structure is fixed, and even when the magnitude of loading is the same in two or three directions.

Results obtained using **dirf** = 0.3 are comparable to the SRSS method (for equal input spectra in each direction), but may be as much as 8% unconservative or 4% over-conservative, depending upon the coordinate system. Larger values of **dirf** tend to produce more conservative results.

Response-Spectrum Analysis Output

Information about each response-spectrum Load Case is available for display, printing, and export using the SAP2000 database tables. This information is described in the following subtopics.

Damping and Accelerations

The modal damping and the ground accelerations acting in each direction are given for every Mode.

The damping value printed for each Mode is the sum of the specified damping for the Load Case, plus the modal damping contributed by effective damping in the Link/Support elements, if any, and the composite modal damping specified in the Material Properties, if any.

The accelerations printed for each Mode are the actual values as interpolated at the modal period from the response-spectrum curves after scaling by the specified value of **sf** and modification for damping. The accelerations are always referred to the local axes of the response-spectrum analysis. They are identified in the output as U1, U2, and U3.

Modal Amplitudes

The response-spectrum modal amplitudes give the multipliers of the mode shapes that contribute to the displaced shape of the structure for each direction of Acceleration. For a given Mode and a given direction of acceleration, this is the product of the modal participation factor and the response-spectrum acceleration, divided by the eigenvalue, ω^2 , of the Mode.

This amplitude, multiplied by any modal response quantity (displacement, force, stress, etc.), gives the contribution of that mode to the value of the same response quantity reported for the response-spectrum load case.

The acceleration directions are always referred to the local axes of the response-spectrum analysis. They are identified in the output as U1, U2, and U3.

For more information:

- See the previous Topic “Damping and Acceleration” for the definition of the response-spectrum accelerations.

- See Topic “Modal Analysis Output” (page 349) in Chapter “Modal Analysis” for the definition of the modal participation factors and the eigenvalues.

Base Reactions

The base reactions are the total forces and moments about the global origin required of the supports (Restrains, Springs, and one-joint Link/Support elements) to resist the inertia forces due to response-spectrum loading.

These are reported separately for each individual Mode and each direction of loading without any combination. The total response-spectrum reactions are then reported after performing modal combination and directional combination.

The reaction forces and moments are always referred to the local axes of the response-spectrum analysis. They are identified in the output as F1, F2, F3, M1, M2, and M3.

Important Note: *Accurate base reactions are best obtained when static-correction modes are included in an eigen analysis, or when Ritz vectors are used. This is particularly true when large stiffnesses are used at the supports and the model is sensitive or ill-conditioned.*

Linear Time-History Analysis

Time-history analysis is a step-by-step analysis of the dynamical response of a structure to a specified loading that may vary with time. The analysis may be linear or nonlinear. This Chapter describes time-history analysis in general, and linear time-history analysis in particular. See Chapter “Nonlinear Time-History Analysis” (page 447) for additional information that applies only to nonlinear time-history analysis.

Basic Topics for All Users

- Overview

Advanced Topics

- Loading
- Initial Conditions
- Time Steps
- Modal Time-History Analysis
- Direct-Integration Time-History Analysis

Overview

Time-history analysis is used to determine the dynamic response of a structure to arbitrary loading. The dynamic equilibrium equations to be solved are given by:

$$\mathbf{K} \mathbf{u}(t) + \mathbf{C} \dot{\mathbf{u}}(t) + \mathbf{M} \ddot{\mathbf{u}}(t) = \mathbf{r}(t)$$

where \mathbf{K} is the stiffness matrix; \mathbf{C} is the damping matrix; \mathbf{M} is the diagonal mass matrix; \mathbf{u} , $\dot{\mathbf{u}}$, and $\ddot{\mathbf{u}}$ are the displacements, velocities, and accelerations of the structure; and \mathbf{r} is the applied load. If the load includes ground acceleration, the displacements, velocities, and accelerations are relative to this ground motion.

Any number of time-history Load Cases can be defined. Each time-history case can differ in the load applied and in the type of analysis to be performed.

There are several options that determine the type of time-history analysis to be performed:

- Linear vs. Nonlinear.
- Modal vs. Direct-integration: These are two different solution methods, each with advantages and disadvantages. Under ideal circumstances, both methods should yield the same results to a given problem.
- Transient vs. Periodic: Transient analysis considers the applied load as a one-time event, with a beginning and end. Periodic analysis considers the load to repeat indefinitely, with all transient response damped out.

Periodic analysis is only available for linear modal time-history analysis.

This Chapter describes linear analysis; nonlinear analysis is described in Chapter “Nonlinear Time-History Analysis” (page 447). However, you should read the present Chapter first.

Loading

The load, $\mathbf{r}(t)$, applied in a given time-history case may be an arbitrary function of space and time. It can be written as a finite sum of spatial load vectors, \mathbf{p}_i , multiplied by time functions, $f_i(t)$, as:

$$\mathbf{r}(t) = \sum_i f_i(t) \mathbf{p}_i \quad (\text{Eqn. 1})$$

The program uses Load Patterns and/or Acceleration Loads to represent the spatial load vectors. The time functions can be arbitrary functions of time or periodic functions such as those produced by wind or sea wave loading.

If Acceleration Loads are used, the displacements, velocities, and accelerations are all measured relative to the ground. The time functions associated with the Acceleration Loads m_x , m_y , and m_z are the corresponding components of uniform ground acceleration, \ddot{u}_{gx} , \ddot{u}_{gy} , and \ddot{u}_{gz} .

Defining the Spatial Load Vectors

To define the spatial load vector, p , for a single term of the loading sum of Equation 1, you may specify either:

- The label of a Load Pattern using the parameter **load**, or
- An Acceleration Load using the parameters **csys**, **ang**, and **acc**, where:
 - **csys** is a fixed coordinate system (the default is zero, indicating the global coordinate system)
 - **ang** is a coordinate angle (the default is zero)
 - **acc** is the Acceleration Load (U1, U2, or U3) in the acceleration local coordinate system as defined below

Each Acceleration Load in the loading sum may have its own acceleration local coordinate system with local axes denoted 1, 2, and 3. The local 3 axis is always the same as the Z axis of coordinate system **csys**. The local 1 and 2 axes coincide with the X and Y axes of **csys** if angle **ang** is zero. Otherwise, **ang** is the angle from the X axis to the local 1 axis, measured counterclockwise when the +Z axis is pointing toward you. This is illustrated in Figure 82 (page 400).

The response-spectrum local axes are always referred to as 1, 2, and 3. The global Acceleration Loads m_x , m_y , and m_z are transformed to the local coordinate system for loading.

It is generally recommended, but not required, that the same coordinate system be used for all Acceleration Loads applied in a given time-history case.

Load Patterns and Acceleration Loads may be mixed in the loading sum.

For more information:

- See Chapter “Load Patterns” (page 321).
- See Topic “Acceleration Loads” (page 338) in Chapter “Load Patterns”.

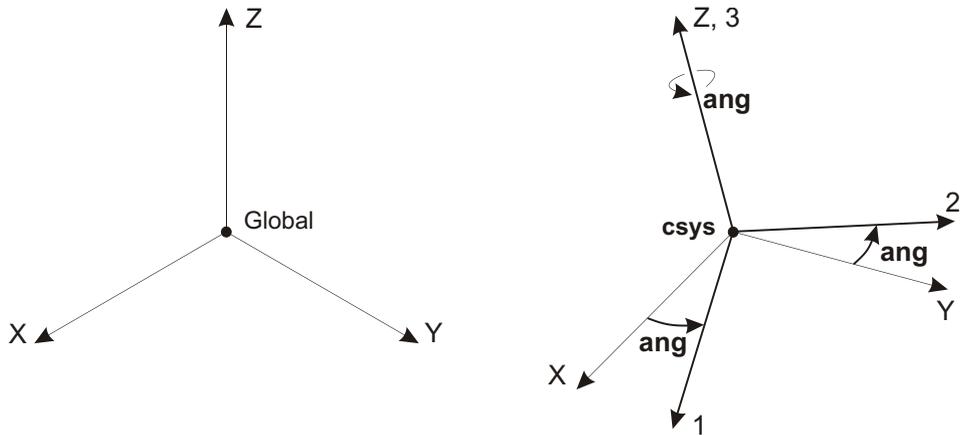


Figure 82
Definition of History Acceleration Local Coordinate System

Defining the Time Functions

To define the time function, $f_i(t)$, for a single term of the loading sum of Equation 1, you may specify:

- The label of a Function, using the parameter **func**, that defines the shape of the time variation (the default is zero, indicating the built-in ramp function defined below)
- A scale factor, **sf**, that multiplies the ordinate values of the Function (the default is unity)
- A time-scale factor, **tf**, that multiplies the time (abscissa) values of the Function (the default is unity)
- An arrival time, **at**, when the Function begins to act on the structure (the default is zero)

The time function, $f_i(t)$, is related to the specified Function, **func(t)**, by:

$$f_i(t) = \mathbf{sf} \cdot \mathbf{func}(t)$$

The analysis time, t , is related to the time scale, **t**, of the specified Function by:

$$t = \mathbf{at} + \mathbf{tf} \cdot \mathbf{t}$$

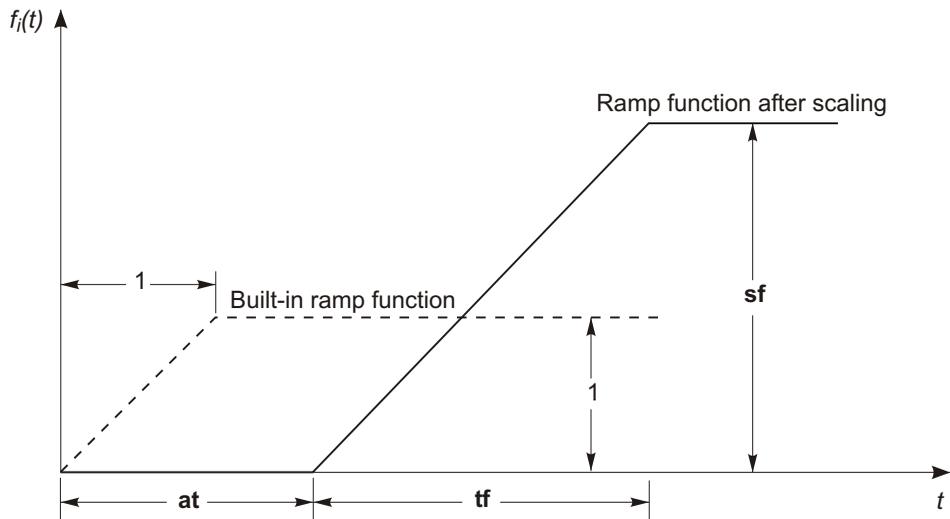


Figure 83
Built-in Ramp Function before and after Scaling

If the arrival time is positive, the application of Function **func** is delayed until after the start of the analysis. If the arrival time is negative, that portion of Function **func** occurring before $t = -at / tf$ is ignored.

For a Function **func** defined from initial time t_0 to final time t_n , the value of the Function for all time $t < t_0$ is taken as zero, and the value of the Function for all time $t > t_n$ is held constant at f_n , the value at t_n .

If no Function is specified, or **func** = 0, the built-in ramp function is used. This function increases linearly from zero at $t = 0$ to unity at $t = 1$ and for all time thereafter. When combined with the scaling parameters, this defines a function that increases linearly from zero at $t = at$ to a value of **sf** at $t = at + tf$ and for all time thereafter, as illustrated in Figure 83 (page 401). This function is most commonly used to gradually apply static loads, but can also be used to build up triangular pulses and more complicated functions.

See Topic “Functions” (page 350) in Chapter “Load Cases” for more information.

Initial Conditions

The initial conditions describe the state of the structure at the beginning of a time-history case. These include:

- Displacements and velocities
- Internal forces and stresses
- Internal state variables for nonlinear elements
- Energy values for the structure
- External loads

The accelerations are not considered initial conditions, but are computed from the equilibrium equation.

For linear transient analyses, zero initial conditions are always assumed.

For periodic analyses, the program automatically adjusts the initial conditions at the start of the analysis to be equal to the conditions at the end of the analysis

If you are using the stiffness from the end of a nonlinear analysis, nonlinear elements (if any) are locked into the state that existed at the end of the nonlinear analysis. For example, suppose you performed a nonlinear analysis of a model containing tension-only frame elements (compression limit set to zero), and used the stiffness from this case for a linear time-history analysis. Elements that were in tension at the end of the nonlinear analysis would have full axial stiffness in the linear time-history analysis, and elements that were in compression at the end of the nonlinear analysis would have zero stiffness. These stiffnesses would be fixed for the duration of the linear time-history analysis, regardless of the direction of loading.

See also Topic “Mass Source” (page 334) in Chapter “Load Patterns” for additional information.

Time Steps

Time-history analysis is performed at discrete time steps. You may specify the number of output time steps with parameter **nstep** and the size of the time steps with parameter **dt**.

The time span over which the analysis is carried out is given by **nstep·dt**. For periodic analysis, the period of the cyclic loading function is assumed to be equal to this time span.

Responses are calculated at the end of each \mathbf{dt} time increment, resulting in $\mathbf{nstep}+1$ values for each output response quantity.

Response is also calculated, but not saved, at every time step of the input time functions in order to accurately capture the full effect of the loading. These time steps are called load steps. For modal time-history analysis, this has little effect on efficiency.

For direct-integration time-history analysis, this may cause the stiffness matrix to be re-solved if the load step size keeps changing. For example, if the output time step is 0.01 and the input time step is 0.005, the program will use a constant internal time-step of 0.005. However, if the input time step is 0.075, then the input and output steps are out of synchrony, and the load steps will be: 0.075, 0.025, 0.05, 0.05, 0.025, 0.075, and so on. For this reason, it is usually advisable to choose an output time step that evenly divides, or is evenly divided by, the input time steps.

Modal Time-History Analysis

Modal superposition provides a highly efficient and accurate procedure for performing time-history analysis. Closed-form integration of the modal equations is used to compute the response, assuming linear variation of the time functions, $f_i(t)$, between the input data time points. Therefore, numerical instability problems are never encountered, and the time increment may be any sampling value that is deemed fine enough to capture the maximum response values. One-tenth of the time period of the highest mode is usually recommended; however, a larger value may give an equally accurate sampling if the contribution of the higher modes is small.

The modes used are computed in a Modal Load Case that you define. They can be the undamped free-vibration Modes (eigenvectors) or the load-dependent Ritz-vector Modes.

If all of the spatial load vectors, p_i , are used as starting load vectors for Ritz-vector analysis, then the Ritz vectors will always produce more accurate results than if the same number of eigenvectors is used. Since the Ritz-vector algorithm is faster than the eigenvector algorithm, the former is recommended for time-history analyses.

It is up to you to determine if the Modes calculated by the program are adequate to represent the time-history response to the applied load. You should check:

- That enough Modes have been computed
- That the Modes cover an adequate frequency range

- That the dynamic load (mass) participation mass ratios are adequate for the Load Patterns and/or Acceleration Loads being applied
- That the modes shapes adequately represent all desired deformations

***Important Note:** Accurate base reactions are best obtained when static-correction modes are included in an eigen analysis, or when Ritz vectors are used. This is particularly true when large stiffnesses are used at the supports and the model is sensitive or ill-conditioned. The use of static-correction modes or Ritz vectors is generally recommended for all modal time-history analyses.*

See Chapter “Modal Analysis” (page 365) for more information.

Modal Damping

The damping in the structure is modeled using uncoupled modal damping. Each mode has a damping ratio, **damp**, which is measured as a fraction of critical damping and must satisfy:

$$0 \leq \text{damp} \leq 1$$

Modal damping has three different sources, which are described in the following. Damping from these sources is added together. The program automatically makes sure that the total is less than one.

Modal Damping from the Load Case

For each linear modal time-history Load Case, you may specify modal damping ratios that are:

- Constant for all modes
- Linearly interpolated by period or frequency. You specify the damping ratio at a series of frequency or period points. Between specified points the damping is linearly interpolated. Outside the specified range, the damping ratio is constant at the value given for the closest specified point.
- Mass and stiffness proportional. This mimics the proportional damping used for direct-integration, except that the damping value is never allowed to exceed unity.

In addition, you may optionally specify damping overwrites. These are specific values of damping to be used for specific modes that replace the damping obtained by one of the methods above. The use of damping overwrites is rarely necessary.

Composite Modal Damping from the Materials

Modal damping ratios, if any, that have been specified for the Materials are converted automatically to composite modal damping. Any cross coupling between the modes is ignored. These modal-damping values will generally be different for each mode, depending upon how much deformation each mode causes in the elements composed of the different Materials.

Effective Damping from the Link/Support Elements

Linear effective-damping coefficients, if any, that have been specified for Link/Support elements in the model are automatically converted to modal damping. Any cross coupling between the modes is ignored. These effective modal-damping values will generally be different for each mode, depending upon how much deformation each mode causes in the Link/Support elements.

Direct-Integration Time-History Analysis

Direct integration of the full equations of motion without the use of modal superposition is available in SAP2000. While modal superposition is usually more accurate and efficient, direct-integration does offer the following advantages for linear problems:

- Full damping that couples the modes can be considered
- Impact and wave propagation problems that might excite a large number of modes may be more efficiently solved by direct integration

For nonlinear problems, direct integration also allows consideration of more types of nonlinearity that does modal superposition.

Direct integration results are extremely sensitive to time-step size in a way that is not true for modal superposition. You should always run your direct-integration analyses with decreasing time-step sizes until the step size is small enough that results are no longer affected by it.

In particular, you should check stiff and localized response quantities. For example, a much smaller time step may be required to get accurate results for the axial force in a stiff member than for the lateral displacement at the top of a structure.

Time Integration Parameters

A variety of common methods are available for performing direct-integration time-history analysis. Since these are well documented in standard textbooks, we will not describe them further here, except to suggest that you use the default “Hilber-Hughes-Taylor alpha” (HHT) method, unless you have a specific preference for a different method.

The HHT method uses a single parameter called alpha. This parameter may take values between 0 and $-1/3$.

For $\alpha = 0$, the method is equivalent to the Newmark method with $\gamma = 0.5$ and $\beta = 0.25$, which is the same as the average acceleration method (also called the trapezoidal rule.) Using $\alpha = 0$ offers the highest accuracy of the available methods, but may permit excessive vibrations in the higher frequency modes, i.e., those modes with periods of the same order as or less than the time-step size.

For more negative values of alpha, the higher frequency modes are more severely damped. This is not physical damping, since it decreases as smaller time-steps are used. However, it is often necessary to use a negative value of alpha to encourage a nonlinear solution to converge.

For best results, use the smallest time step practical, and select alpha as close to zero as possible. Try different values of alpha and time-step size to be sure that the solution is not too dependent upon these parameters.

Damping

In direct-integration time-history analysis, the damping in the structure is modeled using a full damping matrix. Unlike modal damping, this allows coupling between the modes to be considered.

Direct-integration damping has three different sources, which are described in the following. Damping from these sources is added together.

Proportional Damping from the Load Case

For each direct-integration time-history Load Case, you may specify proportional damping coefficients that apply to the structure as a whole. The damping matrix is calculated as a linear combination of the stiffness matrix scaled by a coefficient, c_K , and the mass matrix scaled by a second coefficient, c_M .

You may specify these two coefficients directly, or they may be computed by specifying equivalent fractions of critical modal damping at two different periods or frequencies.

For each nonlinear element in the structure, the coefficient c_K multiplies the initial stiffness matrix, i.e., the stiffness of the element at zero initial conditions, regardless of the nonlinear state used to start this analysis. The exception to this rule is that if the starting nonlinear state has zero stiffness *and* zero force or stress (such as an open gap or a cracked concrete material), then zero damping is assumed. In the case where the initial stiffness is different in the negative and positive direction of loading, the larger stiffness is used.

For cable elements, the damping matrix is proportional to the stiffness matrix for an equivalent truss element along the current chord having the same axial stiffness (AE/L), where L is the undeformed length.

The resulting stiffness-proportional damping is linearly proportional to frequency. It is related to the deformations within the structure. Stiffness proportional damping may excessively damp out high frequency components.

The resulting mass-proportional damping is linearly proportional to period. It is related to the motion of the structure, as if the structure is moving through a viscous fluid. Mass proportional damping may excessively damp out long period components.

Proportional Damping from the Materials

You may specify stiffness and mass proportional damping coefficients for individual materials. For example, you may want to use larger coefficients for soil materials than for steel or concrete. The same interpretation of these coefficients applies as described above for the Load Case damping. Be sure not to double-count the damping by including the same amount in both the Load Case and the materials.

Effective Damping from the Link/Support Elements

Linear effective-damping coefficients, if any, that have been specified for Link/Support elements are directly included in the damping matrix.

Geometric Nonlinearity

SAP2000 is capable of considering geometric nonlinearity in the form of either P-delta effects or large-displacement/rotation effects. Strains within the elements are assumed to be small. Geometric nonlinearity can be considered on a step-by-step basis in nonlinear static and direct-integration time-history analysis, and incorporated in the stiffness matrix for linear analyses.

Advanced Topics

- Overview
- Nonlinear Load Cases
- The P-Delta Effect
- Initial P-Delta Analysis
- Large Displacements

Overview

When the load acting on a structure and the resulting deflections are small enough, the load-deflection relationship for the structure is linear. For the most part, SAP2000 analyses assume such linear behavior. This permits the program to form the equilibrium equations using the original (undeformed) geometry of the struc-

ture. Strictly speaking, the equilibrium equations should actually refer to the geometry of the structure after deformation.

The linear equilibrium equations are independent of the applied load and the resulting deflection. Thus the results of different static and/or dynamic loads can be superposed (scaled and added), resulting in great computational efficiency.

If the load on the structure and/or the resulting deflections are large, then the load-deflection behavior may become nonlinear. Several causes of this nonlinear behavior can be identified:

- **P-delta (large-stress) effect:** when large stresses (or forces and moments) are present within a structure, equilibrium equations written for the original and the deformed geometries may differ significantly, even if the deformations are very small.
- **Large-displacement effect:** when a structure undergoes large deformation (in particular, large strains and rotations), the usual engineering stress and strain measures no longer apply, and the equilibrium equations must be written for the deformed geometry. This is true even if the stresses are small.
- **Material nonlinearity:** when a material is strained beyond its proportional limit, the stress-strain relationship is no longer linear. Plastic materials strained beyond the yield point may exhibit history-dependent behavior. Material nonlinearity may affect the load-deflection behavior of a structure even when the equilibrium equations for the original geometry are still valid.
- **Other effects:** Other sources of nonlinearity are also possible, including nonlinear loads, boundary conditions and constraints.

The large-stress and large-displacement effects are both termed geometric (or kinematic) nonlinearity, as distinguished from material nonlinearity. Kinematic nonlinearity may also be referred to as second-order geometric effects.

This Chapter deals with the geometric nonlinearity effects that can be analyzed using SAP2000. For each nonlinear static and nonlinear direct-integration time-history analysis, you may choose to consider:

- No geometric nonlinearity
- P-delta effects only
- Large displacement and P-delta effects

The large displacement effect in SAP2000 includes only the effects of large translations and rotations. The strains are assumed to be small in all elements.

Material nonlinearity is discussed in Chapters “The Frame Element” (page 105), “Frame Hinge Properties” (page 147), and “The Link/Support Element—Basic” (page 251). Since small strains are assumed, material nonlinearity and geometric nonlinearity effects are independent.

Once a nonlinear analysis has been performed, its final stiffness matrix can be used for subsequent linear analyses. Any geometric nonlinearity considered in the nonlinear analysis will affect the linear results. In particular, this can be used to include relatively constant P-delta effects in buildings or the tension-stiffening effects in cable structures into a series of superposable linear analyses.

For more information:

- See Chapter “Load Cases” (page 341)
- See Chapter “Nonlinear Static Analysis” (page 425)
- See Chapter “Nonlinear Time-History Analysis” (page 447)

Nonlinear Load Cases

For nonlinear static and nonlinear direct-integration time-history analysis, you may choose the type of geometric nonlinearity to consider:

- **None:** All equilibrium equations are considered in the undeformed configuration of the structure
- **P-delta only:** The equilibrium equations take into partial account the deformed configuration of the structure. Tensile forces tend to resist the rotation of elements and stiffen the structure, and compressive forces tend to enhance the rotation of elements and destabilize the structure. This may require a moderate amount of iteration.
- **Large displacements:** All equilibrium equations are written in the deformed configuration of the structure. This may require a large amount of iteration; Newton-Raphson iterations are usually most effective. Although large displacement and large rotation effects are modeled, all strains are assumed to be small. P-delta effects are included.

When continuing one nonlinear Load Case from another, it is recommended that they both have the same geometric-nonlinearity settings.

The large displacement option should be used for any structures undergoing significant deformation; and for buckling analysis, particularly for snap-through buckling and post-buckling behavior. Cables (modeled by frame elements) and other el-

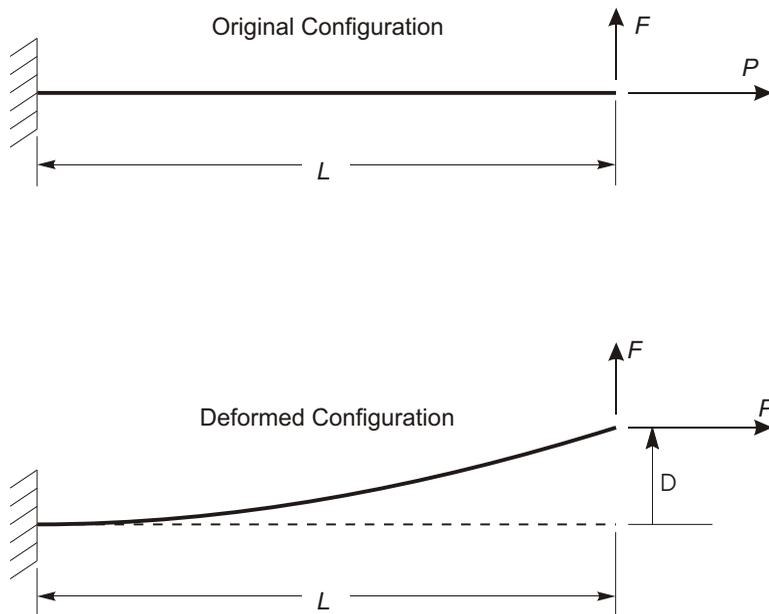


Figure 84
Geometry for Cantilever Beam Example

elements that undergo significant relative rotations within the element should be divided into smaller elements to satisfy the requirement that the strains and relative rotations within an element are small.

For most other structures, the P-delta option is adequate, particularly when material nonlinearity dominates.

If reasonable, it is recommended that the analysis be performed first without geometric nonlinearity, adding P-delta, and possibly large-displacement effects later.

Geometric nonlinearity is not available for nonlinear modal time-history (FNA) analyses, except for the fixed effects that may have been included in the stiffness matrix used to generate the modes.

Note that the catenary Cable element does not require P-delta or Large Displacements to exhibit its internal geometric nonlinearity. The choice should be determined by the rest of the structure.

The P-Delta Effect

The **P-Delta effect** refers specifically to the nonlinear geometric effect of a large tensile or compressive direct stress upon transverse bending and shear behavior. A compressive stress tends to make a structural member more flexible in transverse bending and shear, whereas a tensile stress tends to stiffen the member against transverse deformation.

This option is particularly useful for considering the effect of gravity loads upon the lateral stiffness of building structures, as required by certain design codes (ACI 2002; AISC 2003). It can also be used for the analysis of some cable structures, such as suspension bridges, cable-stayed bridges, and guyed towers. Other applications are possible.

The basic concepts behind the P-Delta effect are illustrated in the following example. Consider a cantilever beam subject to an axial load P and a transverse tip load F as shown in Figure 84 (page 412). The internal axial force throughout the member is also equal to P .

If equilibrium is examined in the original configuration (using the undeformed geometry), the moment at the base is $M = FL$, and decreases linearly to zero at the loaded end. If, instead, equilibrium is considered in the deformed configuration, there is an additional moment caused by the axial force P acting on the transverse tip displacement, δ . The moment no longer varies linearly along the length; the variation depends instead upon the deflected shape. The moment at the base is now $M = FL + P\delta$. The moment diagrams for various cases are shown in Figure 85 (page 414).

Note that only the transverse deflection is considered in the deformed configuration. Any change in moment due to a change in length of the member is neglected here.

If the beam is in tension, the moment at the base and throughout the member is reduced, hence the transverse bending deflection, δ , is also reduced. Thus the member is effectively stiffer against the transverse load F .

Conversely, if the beam is in compression, the moment throughout the member, and hence the transverse bending deflection, δ , are now increased. The member is effectively more flexible against the load F .

If the compressive force is large enough, the transverse stiffness goes to zero and hence the deflection δ tends to infinity; the structure is said to have buckled. The

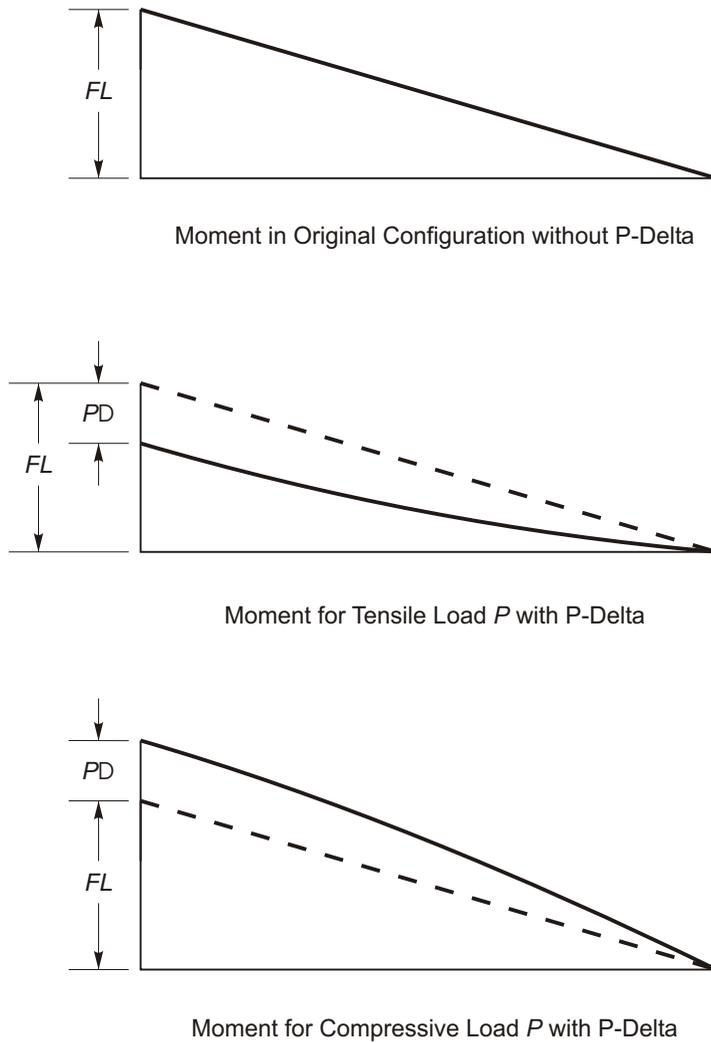


Figure 85
Moment Diagrams for Cantilever Beam Examples

theoretical value of P at which this occurs is called the Euler buckling load for the beam; it is denoted by P_{cr} and is given by the formula

$$P_{cr} = \frac{\pi^2 EI}{4L^2}$$

where EI is the bending stiffness of the beam section.

The exact P-Delta effect of the axial load upon the transverse deflection and stiffness is a rather complicated function of the ratio of the force P to the buckling load P_{cr} . The true deflected shape of the beam, and hence the effect upon the moment diagram, is described by cubic functions under zero axial load, hyperbolic functions under tension, and trigonometric functions under compression.

The P-Delta effect can be present in any other beam configuration, such as simply-supported, fixed-fixed, etc. The P-Delta effect may apply locally to individual members, or globally to the structural system as a whole.

The key feature is that a large axial force, acting upon a small transverse deflection, produces a significant moment that affects the behavior of the member or structure. If the deflection is small, then the moment produced is proportional to the deflection.

P-Delta Forces in the Frame Element

The implementation of the P-Delta effect in the Frame element is described in the following subtopics.

Cubic Deflected Shape

The P-Delta effect is integrated along the length of each Frame element, taking into account the deflection within the element. For this purpose the transverse deflected shape is assumed to be cubic for bending and linear for shear between the rigid ends of the element. The length of the rigid ends is the product of the rigid-end factor and the end offsets, and is usually zero. See Topic “End Offsets” (page 127) in Chapter “The Frame Element” for more information.

The true deflected shape may differ somewhat from this assumed cubic/linear deflection in the following situations:

- The element has non-prismatic Section properties. In this case the P-Delta deflected shape is computed as if the element were prismatic using the average of the properties over the length of the element
- Loads are acting along the length of the element. In this case the P-Delta deflected shape is computed using the equivalent fixed-end forces applied to the ends of the element.

- A large P-force is acting on the element. The true deflected shape is actually described by trigonometric functions under large compression, and by hyperbolic functions under large tension.

The assumed cubic shape is usually a good approximation to these shapes except under a compressive P-force near the buckling load with certain end restraints. Excellent results, however, can be obtained by dividing any structural member into two or more Frame elements. See the *Software Verification Manual* for more detail.

Computed P-Delta Axial Forces

The P-Delta axial force in each Frame element is determined from the axial displacements computed in the element. For meaningful results, it is important to use realistic values for the axial stiffness of these elements. The axial stiffness is determined from the Section properties that define the cross-sectional area and the modulus of elasticity. Using values that are too small may underestimate the P-Delta effect. Using values that are too large may make the P-Delta force in the element very sensitive to the iteration process.

Elements that have an axial force release, or that are constrained against axial deformation by a Constraint, will have a zero P-Delta axial force and hence no P-Delta effect.

The P-Delta axial force also includes loads that act within the element itself. These may include Self-Weight and Gravity Loads, Concentrated and Distributed Span Loads, Prestress Load, and Temperature Load.

The P-Delta axial force is assumed to be constant over the length of each Frame element. If the P-Delta load combination includes loads that cause the axial force to vary, then the average axial force is used for computing the P-Delta effect. If the difference in axial force between the two ends of an element is small compared to the average axial force, then this approximation is usually reasonable. This would normally be the case for the columns in a building structure. If the difference is large, then the element should be divided into many smaller Frame elements wherever the P-Delta effect is important. An example of the latter case could be a flagpole under self-weight.

For more information:

- See Topic “Section Properties” (page 114) in Chapter “The Frame Element.”
- See Topic “End Releases” (page 131) in Chapter “The Frame Element.”
- See Chapter “Constraints and Welds” (page 49).

Prestress

When Prestress Load is included in the P-Delta load combination, the combined tension in the prestressing cables tends to stiffen the Frame elements against transverse deflections. This is true regardless of any axial end releases. Axial compression of the Frame element due to Prestress Load may reduce this stiffening effect, perhaps to zero.

See Topic “Prestress Load” (page 138) in Chapter “The Frame Element” for more information.

Directly Specified P-delta Axial Forces

You may directly specify P-delta forces known to be acting on Frame elements. This is an old-fashioned feature that can be used to model cable structures where the tensions are large and well-known. No iterative analysis is required to include the effect of directly specified P-Delta axial forces.

Use of this feature is not usually recommended! The program does not check if the forces you specify are in equilibrium with any other part of the structure. The directly specified forces apply in *all* analyses and are *in addition to* any P-delta affects calculated in a nonlinear analysis.

We recommend instead that you perform a nonlinear analysis including P-delta or large-displacement effects.

If you use directly specified P-delta forces, you should treat them as if they were a section property that always affects the behavior of the element.

You can assign directly specified P-Delta force to any Frame element using the following parameters:

- The P-Delta axial force, **p**
- A fixed coordinate system, **csys** (the default is zero, indicating the global coordinate system)
- The projection, **px**, of the P-Delta axial force upon the X axis of **csys**
- The projection, **py**, of the P-Delta axial force upon the Y axis of **csys**
- The projection, **pz**, of the P-Delta axial force upon the Z axis of **csys**

Normally only one of the parameters **p**, **px**, **py**, or **pz** should be given for each Frame element. If you do choose to specify more than one value, they are additive:

$$P_0 \quad \mathbf{p} \quad \frac{\mathbf{px}}{c_x} \quad \frac{\mathbf{py}}{c_y} \quad \frac{\mathbf{pz}}{c_z}$$

where P_0 is the P-Delta axial force, and c_x , c_y , and c_z are the cosines of the angles between the local 1 axis of the Frame element and the X, Y, and Z axes of coordinate system **csys**, respectively. To avoid division by zero, you may not specify the projection upon any axis of **csys** that is perpendicular to the local 1 axis of the element.

The use of the P-delta axial force projections is convenient, for example, when specifying the tension in the main cable of a suspension bridge, since the horizontal component of the tension is usually the same for all elements.

It is important when directly specifying P-Delta axial forces that you include all significant forces in the structure. The program does not check for equilibrium of the specified P-Delta axial forces. In a suspension bridge, for example, the cable tension is supported at the anchorages, and it is usually sufficient to consider the P-Delta effect only in the main cable (and possibly the towers). On the other hand, the cable tension in a cable-stayed bridge is taken up by the deck and tower, and it is usually necessary to consider the P-Delta effect in all three components.

P-Delta Forces in the Link/Support Element

P-delta effects can only be considered in a Link/Support element if there is stiffness in the axial (U1) degree of freedom to generate an axial force. A transverse displacement in the U2 or U3 direction creates a moment equal to the axial force (P) times the amount of the deflection (delta).

The total P-delta moment is distributed to the joints as the sum of:

- A pair of equal and opposite shear forces at the two ends that cause a moment due to the length of the element
- A moment at End I
- A moment at End J

The shear forces act in the same direction as the shear displacement (delta), and the moments act about the respectively perpendicular bending axes.

For each direction of shear displacement, you can specify three corresponding fractions that indicate how the total P-delta moment is to be distributed between the three moments above. These fractions must sum to one.

For any element that has zero length, the fraction specified for the shear forces will be ignored, and the remaining two fractions scaled up so that they sum to one. If both of these fractions are zero, they will be set to 0.5.

You must consider the physical characteristics of the device being modeled by a Link/Support element in order to determine what fractions to specify. Long brace or link objects would normally use the shear force. Short stubby isolators would normally use moments only. A friction-pendulum isolator would normally take all the moment on the dish side rather than on the slider side.

Other Elements

For element types other than the Frame and Link/Support, the stresses in the each element are first determined from the displacements computed in the previous iteration. These stresses are then integrated over the element, with respect to the derivatives of the isoparametric shape functions for that element, to compute a standard geometric stiffness matrix that represents the P-delta effect. This is added to the original elastic stiffness matrix of the element. This formulation produces only forces, no moments, at each joint in the element.

Shell elements that are modeling only plate bending will not produce any P-delta effects, since no in-plane stresses will be developed.

Initial P-Delta Analysis

For many applications, it is adequate to consider the P-delta effect on the structure under one set of loads (usually gravity), and to consider all other analyses as linear using the stiffness matrix developed for this one set of P-delta loads. This enables all analysis results to be superposed for the purposes of design.

To do this, define a nonlinear static Load Case that has, at least, the following features:

- Set the name to, say, “PDELTA”
- Start from zero initial conditions
- Apply the Load Patterns that will cause the P-delta effect; often this will be dead load and a fraction of live load
- For geometric nonlinearity, choose P-delta effects

Other parameters include the number of saved steps, the number of iterations allowed per step, and the convergence tolerance. If the P-delta effect is reasonably

small, the default values are adequate. We are not considering staged construction here, although that could be added.

We will refer to this nonlinear static case as the **initial P-delta case**. You can then define or modify other linear Load Cases so that they use the stiffness from case PDELTA:

- Linear static cases
- A modal Load Cases, say called “PDMODES”
- Linear direct-integration time-history cases
- Moving-Load Load Cases

Other linear Load Cases can be defined that are based on the modes from case PDMODES:

- Response-spectrum cases
- Modal time-history cases

Results from all of these cases are superposable, since they are linear and are based upon the same stiffness matrix.

You may also want to define a buckling Load Case that applies the same loads as does case PDELTA, and that starts from zero conditions (not from case PDELTA). The resulting buckling factors will give you an indication of how far from buckling are the loads that cause the P-delta effect.

Below are some additional guidelines regarding practical use of the P-Delta analysis option. See also the *Software Verification Manual* for example problems.

Building Structures

For most building structures, especially tall buildings, the P-Delta effect of most concern occurs in the columns due to gravity load, including dead and live load. The column axial forces are compressive, making the structure more flexible against lateral loads.

Building codes (ACI 2002; AISC 2003) normally recognize two types of P-Delta effects: the first due to the overall sway of the structure and the second due to the deformation of the member between its ends. The former effect is often significant; it can be accounted for fairly accurately by considering the total vertical load at a story level, which is due to gravity loads and is unaffected by any lateral loads. The latter effect is significant only in very slender columns or columns bent in single

curvature (not the usual case); this requires consideration of axial forces in the members due to both gravity and lateral loads.

SAP2000 can analyze both of these P-Delta effects. However, it is recommended that the former effect be accounted for in the SAP2000 analysis, and the latter effect be accounted for in design by using the applicable building-code moment-magnification factors (White and Hajjar 1991). This is how the SAP2000 design processors for steel frames and concrete frames are set up.

The P-Delta effect due to the sway of the structure can be accounted for accurately and efficiently, even if each column is modeled by a single Frame element, by using the factored dead and live loads in the initial P-delta Load Case. The iterative P-Delta analysis should converge rapidly, usually requiring few iterations.

As an example, suppose that the building code requires the following load combinations to be considered for design:

- (1) 1.4 dead load
- (2) 1.2 dead load + 1.6 live load
- (3) 1.2 dead load + 0.5 live load + 1.3 wind load
- (4) 1.2 dead load + 0.5 live load – 1.3 wind load
- (5) 0.9 dead load + 1.3 wind load
- (6) 0.9 dead load + 1.3 wind load

For this case, the P-Delta effect due to overall sway of the structure can usually be accounted for, conservatively, by specifying the load combination in the initial P-delta Load Case to be 1.2 times the dead load plus 0.5 times the live load. This will accurately account for this effect in load combinations 3 and 4 above, and will conservatively account for this effect in load combinations 5 and 6. This P-delta effect is not generally important in load combinations 1 and 2 since there is no lateral load.

The P-Delta effect due to the deformation of the member between its ends can be accurately analyzed only when separate nonlinear Load Cases are run for each load combination above. Six cases would be needed for the example above. Also, at least two Frame elements per column should be used. Again, it is recommended that this effect be accounted for instead by using the SAP2000 design features.

Cable Structures

The P-Delta effect can be a very important contributor to the stiffness of suspension bridges, cable-stayed bridges, and other cable structures. The lateral stiffness of cables is due almost entirely to tension, since they are very flexible when unstressed.

In many cable structures, the tension in the cables is due primarily to gravity load, and it is relatively unaffected by other loads. If this is the case, it is appropriate to define an initial P-delta Load Case that applies a realistic combination of the dead load and live load. It is important to use realistic values for the P-delta load combination, since the lateral stiffness of the cables is approximately proportional to the P-delta axial forces.

P-delta effects are inherent in any nonlinear analysis of Cable elements. P-delta analysis of the whole structure should be considered if you are concerned about compression in the tower, or in the deck of a cable-stayed bridge.

Because convergence tends to be slower for stiffening than softening structures, the nonlinear P-delta analysis may require many iterations. Twenty or more iterations would not be unusual.

Guyed Towers

In guyed towers and similar structures, the cables are under a large tension produced by mechanical methods that shorten the length of the cables. These structures can be analyzed by the same methods discussed above for cabled bridges.

A Strain or Deformation load can be used to produce the requisite shortening. The P-delta load combination should include this load, and may also include other loads that cause significant axial force in the cables, such as gravity and wind loads. Several analyses may be required to determine the magnitude of the length change needed to produce the desired amount of cable tension.

Large Displacements

Large-displacements analysis considers the equilibrium equations in the deformed configuration of the structure. Large displacements and rotations are accounted for, but strains are assumed to be small. This means that if the position or orientation of an element changes, its effect upon the structure is accounted for. However, if the element changes significantly in shape or size, this effect is ignored.

The program tracks the position of the element using an updated Lagrangian formulation. For Frame, Shell, and Link/Support elements, rotational degrees of freedom are updated assuming that the change in rotational displacements between steps is small. This requires that the analysis use smaller steps than might be required for a P-delta analysis. The accuracy of the results of a large-displacement analysis should be checked by re-running the analysis using a smaller step size and comparing the results.

Large displacement analysis is also more sensitive to convergence tolerance than is P-delta analysis. You should always check your results by re-running the analysis using a smaller convergence tolerance and comparing the results.

Applications

Large-displacement analysis is well suited for the analysis of some cable or membrane structures. Cable structures can be modeled with Frame elements, and membrane structures with full Shell elements (you could also use Plane stress elements). Be sure to divide the cable or membrane into sufficiently small elements so that the relative rotations within each element are small.

The catenary Cable element does not require large-displacements analysis. For most structures with cables, P-delta analysis is sufficient unless you expect significant deflection or rotation of the structure supporting or supported by the cables.

Snap-through buckling problems can be considered using large-displacement analysis. For nonlinear static analysis, this usually requires using displacement control of the load application. More realistic solutions can be obtained using nonlinear direct-integration time-history analysis.

Initial Large-Displacement Analysis

The discussion in Topic “Initial P-Delta Analysis” (page 419) in this Chapter applies equally well for an initial large-displacement analysis. Define the initial nonlinear static Load Case in the same way, select large-displacement effects instead of P-delta effects, and make sure the convergence tolerance is small enough. This case can be used as the basis for all subsequent linear analyses.

Nonlinear Static Analysis

Nonlinear static analysis is can be used for a wide variety of purposes, including: to analyze a structure for material and geometric nonlinearity; to form the P-delta stiffness for subsequent linear analyses; to investigate staged (incremental) construction with time-dependent material behavior; to perform cable analysis; to perform static pushover analysis; and more.

Although much of this Chapter is advanced, basic knowledge of nonlinear static analysis is essential for P-delta analysis and modeling of tension-only braces and cables.

Basic Topics for All Users

- Overview
- Nonlinearity
- Important Considerations
- Loading
- Initial Conditions
- Output Steps

Advanced Topics

- Load Application Control
- Staged Construction
- Nonlinear Solution Control
- Static Pushover Analysis

Overview

Nonlinear static analysis can be used for many purposes:

- To perform an initial P-delta or large-displacement analysis to get the stiffness used for subsequent superposable linear analyses
- To perform staged (incremental, segmental) construction analysis, including material time-dependent effects like aging, creep and shrinkage
- To change the Mass Source to be used for subsequent analyses
- To analyze structures with tension-only bracing
- To analyze cable structures
- To perform static pushover analysis
- To perform snap-through buckling analyses
- To establish the initial conditions for nonlinear direct-integration time-history analyses
- For any other static analysis that considers the effect of material or geometric nonlinear behavior

Any number of nonlinear Static Load Cases can be defined. Each case can include one or more of the features above. In a nonlinear analysis, the stiffness and load may all depend upon the displacements. This requires an iterative solution to the equations of equilibrium.

Nonlinearity

The following types of nonlinearity are available in SAP2000:

- Material nonlinearity
 - Various type of nonlinear properties in Link/Support elements
 - Tension and/or compression limits in Frame elements

- Plastic hinges in Frame elements
- Nonlinear material behavior in layered Shell elements
- Geometric nonlinearity
 - P-delta effects
 - Large displacement effects
- Staged construction
 - Changes in the structure
 - Aging, creep, and shrinkage

All material nonlinearity that has been defined in the model will be considered in a nonlinear static Load Case.

You have a choice of the type of geometric nonlinearity to be considered:

- None
- P-delta effects
- Large displacement effects

If you are continuing from a previous nonlinear analysis, it is strongly recommended that you select the same geometric nonlinearity parameters for the current case as for the previous case. See Chapter “Geometric Nonlinearity” (page 409) for more information.

Staged construction is available as an option. Even if the individual stages are linear, the fact that the structure changes from one stage to the next is considered to be a type of nonlinearity.

Important Considerations

Nonlinear analysis takes time and patience. Each nonlinear problem is different. You can expect to need a certain amount of time to learn the best way to approach each new problem.

Start with a simple model and build up gradually. Make sure the model performs as expected under linear static loads and modal analysis. Rather than starting with nonlinear properties everywhere, add them in increments beginning with the areas where you expect the most nonlinearity.

If you are using frame hinges, start with models that do not lose strength for primary members; you can modify the hinge models later or redesign the structure.

When possible, perform your initial analyses without geometric nonlinearity. Add P-delta effects, and possibly large deformations, much later. Start with modest target displacements and a limited number of steps. In the beginning, the goal should be to perform the analyses quickly so that you can gain experience with your model. As your confidence grows with a particular model you can push it further and consider more extreme nonlinear behavior.

Mathematically, nonlinear static analysis does not always guarantee a unique solution. Inertial effects in dynamic analysis and in the real world limit the path a structure can follow. But this is not true for static analysis, particularly in unstable cases where strength is lost due to material or geometric nonlinearity. If a nonlinear static analysis continues to cause difficulties, change it to a direct-integration time-history analysis and apply the load quasi-statically (very slowly.)

Small changes in properties or loading can cause large changes in nonlinear response. For this reason, it is extremely important that you consider many different loading cases, and that you perform sensitivity studies on the effect of varying the properties of the structure.

Loading

You may apply any combination of Load Patterns, Acceleration Loads, and modal loads.

A modal load is a specialized type of loading used for pushover analysis. It is a pattern of forces on the joints that is proportional to the product of a specified mode shape times its circular frequency squared (ω^2) times the mass tributary to the joint.

The specified combination of loads is applied simultaneously. Normally the loads are applied incrementally from zero to the full specified magnitude. For specialized purposes (e.g., pushover or snap-through buckling), you have the option to control the loading by monitoring a resulting displacement in the structure. See Topic “Load Application Control” (page 428) in this Chapter for more information.

Load Application Control

You may choose between a load-controlled or displacement-controlled nonlinear static analysis. For both options, the pattern of loads acting on the structure is determined by the specified combination of loads. Only the scaling is different.

Normally you would choose load control. It is the most common, physical situation.

Displacement control is an advanced feature for specialized purposes.

Load Control

Select load control when you know the magnitude of load that will be applied and you expect the structure to be able to support that load. An example would be when applying gravity load, since it is governed by nature.

Under load control, all loads are applied incrementally from zero to the full specified magnitude.

Displacement Control

Select displacement control when you know how far you want the structure to move, but you don't know how much load is required. This is most useful for structures that become unstable and may lose load-carrying capacity during the course of the analysis. Typical applications include static pushover analysis and snap-through buckling analysis.

To use displacement control, you must select a displacement component to monitor. This may be a single degree of freedom at a joint, or a generalized displacement that you have previously defined. See Topic "Generalized Displacement" (page 45) in Chapter "Joints and Degrees of Freedom" for more information.

You must also give the magnitude of the displacement that is your target for the analysis. The program will attempt to apply the load to reach that displacement. The load magnitude may be increased and decreased during the analysis.

Be sure to choose a displacement component that monotonically increases during loading. If this is not possible, you may need to divide the analysis into two or more sequential cases, changing the monitored displacement in the different cases. The use of the conjugate displacement control, described below, may help solve this problem for you.

Important note: Using displacement control is NOT the same thing as applying displacement loading on the structure! Displacement control is simply used to MEASURE the displacement at one point that results from the applied loads, and to adjust the magnitude of the loading in an attempt to reach a certain measured displacement value. The overall displaced shape of the structure will be different for different patterns of loading, even if the same displacement is controlled.

Conjugate Displacement Control

If analysis is having trouble reaching the desired displacement target, you can try the option for the load case to use the conjugate displacement for control. The conjugate displacement is a weighted average of all displacements in the structure, each displacement degree of freedom being weighted by the load acting on that degree of freedom. In other words, it is a measure of the work done by the applied load.

This option is particularly useful for unsymmetrical buildings where the top story level may rotate about the vertical axis while being pushed under lateral load. The rotation could cause the monitored displacement to move backward even while the weighted average displacement is moving forward.

If you choose to use the conjugate displacement control, it will be used to determine whether the load should be increased or decreased. The specified monitored displacement will still be used to set the target displacement, i.e., how far the structure should move. However, this target may not be matched exactly.

Initial Conditions

The initial conditions describe the state of the structure at the beginning of a Load Case. These include:

- Displacements and velocities
- Internal forces and stresses
- Internal state variables for nonlinear elements
- Energy values for the structure
- External loads

For a static analysis, the velocities are always taken to be zero.

For nonlinear analyses, you may specify the initial conditions at the start of the analysis. You have two choices:

- Zero initial conditions: the structure has zero displacement and velocity, all elements are unstressed, and there is no history of nonlinear deformation.
- Continue from a previous nonlinear analysis: the displacements, velocities, stresses, loads, energies, and nonlinear state histories from the end of a previous analysis are carried forward.

Nonlinear static and nonlinear direct-integration time-history cases may be chained together in any combination, i.e., both types of analysis are compatible with each other. It is strongly recommended that you select the same geometric nonlinearity parameters for the current case as for the previous case.

When continuing from a previous case, all applied loads specified for the present Load Case are incremental, i.e., they are added to the loads already acting at the end of the previous case.

When multiple Mass Sources have been defined, you can specify which Mass Source to use for a nonlinear static or nonlinear direct-integration time-history load case. However, the default is to use the same Mass Source as the previous load case, or the default Mass Source if starting from zero initial conditions. See Topic “Mass Source” (page 334) in Chapter “Load Patterns” for more information.

Nonlinear static cases cannot be chained together with nonlinear modal time-history (FNA) cases.

Output Steps

Normally, only the final state is saved for a nonlinear static analysis. This is the result after the full load has been applied.

You can choose instead to save intermediate results to see how the structure responded during loading. This is particularly important for static pushover analysis, where you need to develop the capacity curve.

If you are only interested in the saving the final result, you can skip the rest of this topic.

Saving Multiple Steps

If you choose to save multiple states, the state at the beginning of the analysis (step 0) will be saved, as well as a number of intermediate states. From a terminology point of view, saving five steps means the same thing as saving six states (steps 0 to 5): the step is the increment, and the state is the result.

The number of saved steps is determined by the parameters:

- Minimum number of saved steps
- Maximum number of saved steps
- Option to save positive increments only

These are described in the following.

Minimum and Maximum Saved Steps

The Minimum Number of Saved Steps and Maximum Number of Saved Steps provide control over the number of points actually saved in the analysis. If the minimum number of steps saved is too small, you may not have enough points to adequately represent a pushover curve. If the minimum and maximum number of saved steps is too large, then the analysis may consume a considerable amount of disk space, and it may take an excessive amount of time to display results.

The program automatically determines the spacing of steps to be saved as follows. The maximum step length is equal to total force goal or total displacement goal divided by the specified Minimum Number of Saved Steps. The program starts by saving steps at this increment. If a significant event occurs at a step length less than this increment, then the program will save that step too and continue with the maximum increment from there. For example, suppose the Minimum Number of Saved Steps and Maximum Number of Saved Steps are set at 20 and 30 respectively, and the target is to be to a displacement of 10 inches. The maximum increment of saved steps will be $10 / 20 = 0.5$ inches. Thus, data is saved at 0.5, 1, 1.5, 2, 2.5 inches. Suppose that a significant event occurs at 2.7 inches. Then data is also saved at 2.7 inches, and continues on from there being saved at 3.2, 3.7, 4.2, 4.7, 5.2, 5.7, 6.2, 6.7, 7.2, 7.7, 8.2, 8.7, 9.2, 9.7 and 10.0 inches.

The Maximum Number of Saved Steps controls the number of significant events for which data will be saved. The program will always reach the force or displacement goal within the specified number of maximum saved steps, however, in doing so it could have to skip saving steps at later events. For example, suppose the Minimum Saved Steps is set to 20, the Maximum Number of Saved Steps is set to 21, and the pushover is to be to a displacement of 10 inches. The maximum increment of saved steps will be $10 / 20 = 0.5$ inches. Thus, data is saved at 0.5, 1, 1.5, 2, 2.5 inches. Suppose that a significant event occurs at 2.7 inches. Then data is also saved at 2.7 inches, and continues on from there being saved at 3.2 and 3.7 inches. Suppose another significant event occurs at 3.9 inches. The program will not save the data at 3.9 inches because if it did it would not be able to limit the maximum increment to 0.5 inches and still get through the full pushover in no more than 21 steps. Note that if a second significant event occurred at 4.1 inches rather than 3.9 inches, then the program would be able to save the step and still meet the specified criteria for maximum increment and maximum number of steps.

Save Positive Increments Only

This option is primarily of interest for pushover analysis under displacement control. In the case of extreme nonlinearity, particularly when a frame hinge sheds load, the pushover curve may show negative increments in the monitored displacement while the structure is trying to redistribute the force from a failing component.

You may choose whether or not you want to save only the steps having positive increments. The negative increments often make the pushover curve look confusing. However, seeing them can provide insight into the performance of the analysis and the structure.

You may want to choose to Save Positive Increments Only in most cases except when the analysis is having trouble converging.

Nonlinear Solution Control

The specified combination of applied loads is applied incrementally, using as many steps as necessary to satisfy equilibrium and to produce the requested number of saved output steps.

The nonlinear equations are solved iteratively in each load step. This may require re-forming and re-solving the stiffness matrix. The iterations are carried out until the solution converges. If convergence cannot be achieved, the program divides the step into smaller substeps and tries again.

Several parameters are available for you to control the iteration and substepping process. These are described in the following. We recommend that you use the default values of these parameters to start, except that you may often need to increase the maximum number of total steps and null steps for more complex models.

If you are having convergence difficulties, you may try varying the iteration control parameters below. However, you should also consider that the model itself may need improvement. Look for instabilities due to inadequate support, buckling, and excessively large stiffnesses. If you have hinges or materials that lose strength, make sure that this behavior is really necessary and that the negative slopes are not unrealistically too steep.

Maximum Total Steps

This is the maximum number of steps allowed in the analysis. It may include saved steps as well as intermediate substeps whose results are not saved. The purpose of setting this value is to give you control over how long the analysis will run.

Start with a smaller value to get a feel for the time the analysis will take. If an analysis does not reach its target load or displacement before reaching the maximum number of steps, you can re-run the analysis after increasing this maximum number of saved steps. The length of time it takes to run a nonlinear static analysis is approximately proportional to the total number of steps.

Maximum Null (Zero) Steps

Null (zero) steps may occur during the nonlinear solution procedure when:

- Iteration does not converge and a smaller step size is attempted
- Snap-back is detected due to strength loss in a frame hinge

An excessive number of null steps may indicate that the solution is stalled due to instability or numerical sensitivity.

You can set the Maximum Null (Zero) Steps so that the solution will terminate early if it is having trouble converging. Set this value equal to the Maximum Total Steps if you do not want the analysis to terminate due to null steps.

Event-to-Event Stepping Control

Event-to-event stepping is available as an option for nonlinear static analysis. Using this strategy, an increment of load is applied until one of the nonlinear elements experiences an “event” that causes a significant change in the stiffness of the structure. Iteration is performed for this load increment to ensure that equilibrium is achieved to within a specified convergence tolerance. The stiffness matrix is then reformed, and another load increment is applied to reach the next event. This continues in a series of event steps until the entire load has been applied.

Events are implemented for the following nonlinear models: all frame hinges; layered shells with directional or coupled nonlinear behavior; and links with multi-linear plasticity, gap, hook, friction-pendulum, and triple-pendulum behavior. Events typically include first yield, unloading, strength loss, increasing stiffness, and other significant changes in behavior that depend on the type of element or material.

You can specify whether or not to use event-to-event stepping. When events are used, each load increment tends to be more linear, thus reducing equilibrium unbalance and lessening the need for iteration. On the other hand, models with a large number of nonlinear elements and/or hinges may generate a large number of events, which can significantly increase the analysis time.

When using event stepping, you can specify a relative event-lumping tolerance that will group multiple events together. Effectively, the load increment calculated to reach the first event is increased by the tolerance to include other events that would occur soon afterward. Larger event tolerances reduce the number of event steps taken, but may increase the amount of iteration required.

Note that for isotropic frame hinges the event lumping tolerance may also allow the force or deformation response to deviate from the backbone curve. The relative magnitude of the deviations is on the order of the relative lumping tolerance.

Whether or not to use event stepping is a problem-dependent decision. Normally using events will improve convergence behavior. However, if a large number of event steps are generated that make the analysis run slowly, try turning event-stepping off for that model or load case. Either way, results for a stable model should be consistent to within the convergence tolerance.

Nonlinear iteration is always used for nonlinear static analysis, whether or not event-stepping is used. This is different from nonlinear direct-integration time-history analysis where events can be used with or without iteration. Note that the line-search iteration option is not available when event stepping is used for nonlinear static analysis.

Nonlinear Iteration

Iteration is used to ensure that equilibrium is achieved at each step and substep of the analysis to within a specified convergence tolerance. You can specify the relative convergence tolerance used to compare the magnitude of force error with the magnitude of the force acting on the structure.

You may need to use significantly smaller values of convergence tolerance to get good results for large-displacements problems than for other types of nonlinearity. Try decreasing this value until you get consistent results.

For each step, constant-stiffness iteration is tried first. If convergence is not achieved, Newton-Raphson (tangent-stiffness) iteration is tried next. If both fail, the step size is reduced, and the process is repeated.

You can separately control the number of constant-stiffness and Newton-Raphson iterations allowed in each step. Setting either parameter to zero prevents that type of iteration. Setting both to zero causes the analysis to use the default number of iterations. Constant-stiffness iterations are faster than Newton-Raphson iterations, but the latter are usually more effective, especially for cables and geometric nonlinearity. The default values work well in many situations.

Line Search Option

Line search is an iteration option that can increase the efficiency of equilibrium iterations by scaling the solution increment in a trial-and-error fashion to find the smallest unbalance. This increases the computation time for each iteration, but often results in fewer iterations and better convergence behavior, with a net gain in efficiency. Line search is particularly effective for stiffening systems, such as tightening cables, closing gaps, unloading from plastic yielding, and sticking after frictional sliding.

You can choose whether or not to use line search, except that line search is not available when using event-to-event stepping. Line search is also not used for displacement-controlled load cases.

When using line search, specify the following parameters:

- **Maximum number of line searches per Iteration.** Use this parameter to limit the number of line searches allowed per iteration, usually in the range of 5 to 40.
- **Line-Search Acceptance Tolerance (Relative).** The solution increment is successively reduced or increased until the minimum error is found, the maximum number of line searches is reached, or the relative change in error from the previous trial is less than the acceptance tolerance set here. The practical range is about 0.02 to 0.50. This value should not be too small, since the goal is only to improve the next iteration.
- **Line-Search Step Factor.** For each line-search trial, the solution increment is successively reduced by the step factor specified here until the minimum error is found, the maximum number of line searches is reached, or the acceptance tolerance is satisfied. If reducing the increment does not decrease the error, the increment is instead increased by the step factor following the same procedure. This factor must be greater than 1.0, and should usually be no more than 2.0.

The default values are recommended as starting point. For stiffening systems that are slow to converge, try increasing the maximum number of line searches per iteration, decreasing the acceptance tolerance, and reducing the step factor.

Static Pushover Analysis

Nonlinear static pushover analysis is a specialized procedure used in performance-based design for seismic loading. SAP2000 provides the following tools needed for pushover analysis:

- Material nonlinearity at discrete, user-defined hinges in Frame elements. The hinge properties were created with pushover analysis in mind. Default hinge properties are provided based on ASCE 41 and other code-based criteria. See Chapter “Frame Hinge Properties” (page 147). In ETABS, discrete fiber P-M3 hinges can be assigned to wall elements.
- Nonlinear static analysis procedures specially designed to handle the sharp drop-off in load carrying capacity typical of frame hinges used in pushover analysis.
- Nonlinear static analysis procedures that allow displacement control, so that unstable structures can be pushed to desired displacement targets. See Topic “Load Application Control” (page 428) in this Chapter.
- Display capabilities in the graphical user interface to generate and plot pushover curves, including demand and capacity curves in spectral ordinates. See the online Help facility in the graphical user interface for more information.
- Capabilities in the graphical user interface to plot and output the state of every hinge at each step in the pushover analysis. See Chapter “Frame Hinge Properties” (page 147) and the online Help facility in the graphical user interface for more information.

In addition to these specialized features, the full nonlinearity of the program can be used, including nonlinear Shell and Link/Support behavior, geometric nonlinearity, and staged construction. In addition, you are not restricted to static pushover analysis: you can also perform full nonlinear time-history analysis.

The following general sequence of steps is involved in performing nonlinear static pushover analysis using SAP2000:

1. Create a model just like you would for any other analysis.
2. Define frame hinge properties and assign them to the frame elements.

3. Define any Load Patterns and static and dynamic Load Cases that may be needed for steel or concrete design of the frame elements, particularly if default hinges are used.
4. Run the Load Cases needed for design.
5. If any concrete hinge properties are based on default values to be computed by the program, you must perform concrete design so that reinforcing steel is determined.
6. If any steel hinge properties are based on default values to be computed by the program for Auto-Select frame section properties, you must perform steel design and accept the sections chosen by the program.
7. Define the Load Patterns that are needed for use in the pushover analysis, including:
 - Gravity loads and other loads that may be acting on the structure before the lateral seismic loads are applied. You may have already defined these Load Patterns above for design.
 - Lateral loads that will be used to push the structure. If you are going to use Acceleration Loads or modal loads, you don't need any new Load Patterns, although modal loads require you to define a Modal Load Case.
8. Define the nonlinear static Load Cases to be used for pushover analysis, including:
 - A sequence of one or more cases that start from zero and apply gravity and other fixed loads using load control. These cases can include staged construction and geometric nonlinearity.
 - One or more pushover cases that start from this sequence and apply lateral pushover loads. These loads should be applied under displacement control. The monitored displacement is usually at the top of the structure and will be used to plot the pushover curve.
9. Run the pushover Load Cases.
10. Review the pushover results: Plot the pushover curve, the deflected shape showing the hinge states, force and moment plots, and print or display any other results you need.
11. Revise the model as necessary and repeat.

It is important that you consider several different lateral pushover cases to represent different sequences of response that could occur during dynamic loading. In particular, you should push the structure in both the X and Y directions, and possibly at angles in between. For non-symmetrical structures, pushing in the positive and negative direction may yield different results. When pushing in a given direction, you may want to consider different vertical distributions of the lateral load, such as the first and second mode in that direction.

Staged Construction

Staged construction is a special type of nonlinear static analysis that requires a separate add-on module for this feature to become available in the program.

Staged construction allows you to define a sequence of stages wherein you can add or remove portions of the structure, selectively apply load to portions of the structure, and to consider time-dependent material behavior such as aging, creep, and shrinkage. Staged construction is variously known as incremental construction, sequential construction, or segmental construction.

Normally the program analyzes the whole structure in all Load Cases. If you do not want to perform staged-construction analysis, you can skip the rest of this topic.

Staged construction is considered a type of nonlinear static analysis because the structure may change during the course of the analysis. However, consideration of material and geometric nonlinearity is optional. Because staged construction is a type of nonlinear static analysis, it may be part of a sequence of other nonlinear static and direct-integration time-history Load Cases, and it may also be used as a stiffness basis for linear Load Cases.

If you continue any nonlinear analysis from a staged construction analysis, or perform a linear analysis using its stiffness, only the structure as built at the end of the staged construction will be used.

Stages

For each nonlinear staged-construction Load Case, you define a sequence of stages. These are analyzed in the order defined. You can specify as many stages as you want in a single Load Case. Staged construction can also continue from one Load Case to another.

For each stage you specify the following:

- A duration, in days. This is used for time-dependent effects. If you do not want to consider time-dependent effects in a given stage, set the duration to zero.
- Any number of objects to be added to the structure, or none. The age of the objects at the time they are added can be specified, if time-dependent effects are to be considered.
- Any number of objects to be removed from the structure, or none
- Any number of objects to be loaded by specified Load Patterns, or none. You may specify that all objects in the group are to be loaded, or only those objects in the group that are being added to the structure in this stage.
- Property-Modifier Named Sets may be assigned to any number of frame or shell objects. See Chapters “The Frame Element” and “The Shell Element” for more information.
- End-Release Named Sets may be assigned to any number of frame objects. See Chapter “The Frame Element” for more information.
- Section properties may be changed for any number of frame, shell, tendon, and link/support objects.
- Bridge superstructure girders may be flagged as composite or non-composite for design purposes. By default, they are non-composite. This option is available only in CSiBridge, and only for bridge groups that contain top-slab components. This flag is only used in bridge superstructure design checks for constructability. ***It has absolutely no effect upon analysis.***

Objects may be specified individually or by using groups, including specialized groups such as ETABS building stories or CSiBridge superstructure segments. In most cases, the use of groups is convenient, so the first step to setting up staged-construction analysis is to define groups for that purpose. See Topic “Groups” (page 9) in Chapter “Objects and Elements” for more information. Note that there is always a built-in group called "ALL" that includes the whole structure. When an operation only pertains to some object types, only objects of those types in a group will be affected.

When you specify staged construction, the analysis starts with the structure as built from the previous Load Case. If you are starting from zero, then the structure starts with no objects.

Each stage is analyzed separately in the order the stages are defined. The analysis of a stage has two parts:

- (1) Changes to the structure and application of loads are analyzed. These occur instantaneously in time, i.e., the analysis may be incremental, but no time elapses from the point-of-view of the material
- (2) If non-zero duration has been specified, time-dependent material effects are then analyzed. During this time, the structure does not change and applied loads are held constant. However, internal stress redistribution may occur.

The instantaneous part (1) of the stage is analyzed as follows (reference to groups includes individual objects as well):

- The objects to be added, if any, are processed. Only new objects in the specified groups (not already present in the structure) are added. For each non-joint object added, all joints connected to that object are also added, even if they are not explicitly included in the group.
- The groups to be removed, if any, are processed. Only objects actually present in the structure are removed. When objects are removed, their stiffness, mass, loads, and internal forces are removed from the structure and replaced with equivalent forces. These forces are linearly reduced to zero during the course of the analysis. Joints that were automatically added will be removed when all connected objects are removed.
- Section properties changes, if any, are processed as described below.
- Named set assignments, if any, are processed.
- All specified loads will be increased linearly during the course of the analysis. Loads specified on *all* objects in a group will only be applied to objects that are actually present in the structure or are being added in this stage. Loads specified on *added* objects in a group will only be applied to objects that are being added in this stage.

If an object is included in more than one group that is being added or removed, the object will only be added or removed once. Whether it is added or removed depends on which operation occurs last in the order you have specified them. For example, if an object is included in three groups that are being added and one group that is being removed, the object will be removed if that was the last operation specified for that stage. Changing section properties acts as a removal and addition, and it will occur in sequence with explicit addition and removal depending on the order of operations you specify. See below for more information.

If an object is included in more than one group that is being loaded, the object will be multiply loaded.

Load application must be by load control. Displacement control is not allowed.

Changing Section Properties

The main purpose for changing section properties is to be able to model multiple configurations of the same structure. This can be accomplished changing the properties of newly added, unstressed objects. You may, however, change section properties for stressed objects, and it is important that you understand exactly what the analysis is doing so that you can achieve the results you want.

Section property changes on each object are processed as follows:

- If the object is present, it is removed and replaced with equivalent forces so that the structure does not deflect.
- The object is added again, this time with the new section property, but unstressed and without any loading that may have been acting on the object when it was removed. The object will be added even if it was not previously present in the structure.
- The equivalent forces, if any, are then reversed. This will usually deform the object and cause the structure will deflect.

You would normally want to reapply self-weight to any object whose section properties are changed, since previous loading is discarded.

This operation is essentially removing a stressed object and replacing it with an unstressed object. You can apply forces to unload the object before changing its section if you want to simulate removing an unstressed object.

Changing section properties is useful for analysis purposes. However, the object will be treated as if it has its original section property for all display, output, and design purposes.

Output Steps

The specification of output steps is similar to that described earlier in this Chapter in Topic “Output Steps” (page 431), except that you can individually control the number of steps to be saved for the two parts of each stage:

- (1) How many steps to save during changes to the structure and instantaneous application of load
- (2) How many steps to save during the time-dependent analysis for aging, creep and shrinkage.

The number of steps requested for these two parts of each stage applies equally to all stages in the Load Case.

Important! The time step used for the time-dependent analysis is based on the number of steps saved. For statically indeterminate structures where significant stress redistribution may occur due to creep and shrinkage, it is important to use small-enough time steps, especially during the youth of the structure where large changes may be occurring. You may want to re-run the analysis with increasing numbers of steps saved until you are satisfied that the results have converged.

Example

Let's build a simple bridge. Define four groups: "BENTS," "DECK1," "DECK2," "SHORING," and "APPURTS." The structure can be linear or nonlinear. Time-dependent properties are assumed for the concrete material.

Also define three Load Patterns:

- "GRAVITY", to apply dead load
- "TENSION", to apply post-tensioning cable loads
- "EQUIPMENT", to apply the weight of temporary construction equipment on the deck

Define a staged-construction Load Case called "BUILD" that starts from zero, and contains the following stages:

1. a) Add group "BENTS" with an age of 10 days.
b) Apply load "GRAVITY" to added elements in group "ALL".
c) No time-dependent effects need to be considered—we can assume these already happened in the first 10 days.
2. a) Add group "SHORING" with an age of 10 days.
b) Apply load "GRAVITY" to added elements in group "ALL".
c) No time-dependent effects need to be considered for the shoring.
3. a) Add group "DECK1" with an age of 0 days (wet concrete).
b) Apply load "GRAVITY" to added elements in group "ALL".
c) Apply load "EQUIPMENT" to all elements in group "DECK1".
d) Allow 3 days duration for aging, creep, and shrinkage.
4. a) Add group "DECK2" with an age of 0 days (wet concrete).
b) Apply load "GRAVITY" to added elements in group "ALL".
c) Remove load "EQUIPMENT" from all elements in group "DECK1" (apply

- with a scale factor of -1.0).
- d) Apply load “EQUIPMENT” to all elements in group “DECK2”.
 - e) Allow 3 days duration for aging, creep, and shrinkage.
5. a) Remove load “EQUIPMENT” from all elements in group “DECK2” (apply with a scale factor of -1.0).
- b) Apply load “TENSION” to all elements in group “ALL”.
4. a) Remove group “SHORING”.
- b) Allow 3 days duration for aging, creep, and shrinkage.
5. a) Add group “APPURTS” with an age of 10 days.
- b) Allow 30 days duration for aging, creep, and shrinkage.
6. a) Allow 300 days duration for aging, creep, and shrinkage.
7. a) Allow 3000 days duration for aging, creep, and shrinkage.

The reason for adding several stages with increasing length of time at the end is to get long term effects at increasing time-step size, since the number of output steps is the same for all stages.

Case BUILD can now be used to define the stiffness matrix for any number of linear analyses, including modal, response-spectrum, moving-load, and other types. You can also continue case BUILD with a nonlinear direct-integration time-history analysis for seismic load, or even more nonlinear static cases that may include pushover analysis or more staged construction for the purposes of retrofit.

Target-Force Iteration

When any Load Pattern containing target-force loads is applied in a nonlinear static Load Case, internal deformation load is iteratively applied to achieve the target force. In a staged-construction analysis, the iteration occurs individually over any stage for which target-force loads are applied. Otherwise, the iteration is for the whole nonlinear static Load Case.

Trial deformation loads are applied to those elements for which target forces have been assigned, and a complete nonlinear analysis (or stage thereof) is performed. At the end of the analysis, and assuming that convergence for equilibrium has been achieved, forces in the targeted elements are compared with their desired targets. A relative error is computed that is the root-mean-square over all the elements of the difference between the target and the actual force, divided by the larger of the two values. If this error is greater than the relative convergence tolerance that you speci-

fied, a revised deformation load is computed and the complete nonlinear analysis (or stage thereof) is performed again. This process is repeated until the error is less than the tolerance, or the specified maximum number of iterations is reached.

As part of the definition of the nonlinear static Load Case, you may specify the following parameters to control target-force iteration:

- **Relative convergence tolerance:** This is the error you are willing to accept in the target forces. Since target forces represent your desire, and not a natural requirement like equilibrium, a large value such as 0.01 to 0.10 is recommended.
- **Maximum number of iterations:** Many iterations may be required for some structures, and each iteration may take some time. Start with a moderate value, such as 5 to 10, and increase it as necessary.
- **Acceleration factor:** The difference between the target force and the force actually achieved is converted to deformation load and applied in the next iteration. You can increase or decrease this load by the acceleration factor. Use a value greater than one if convergence is slow, such as when pulling or pushing against a flexible structure. Use a value less than one if the solution is diverging, i.e., the unbalance is growing or oscillating between iterations.
- **Whether to continue the analysis if convergence is not achieved:** Since achieving the target forces is not a natural requirement, and may not even be possible, you may want to continue the analysis even if the target is not reached.

You should be realistic in your expectations for target-force iteration. You cannot arbitrarily specify the forces in a statically determinate structure, such as a truss. Convergence will be slow when target forces are specified in elements connected to very flexible supports, or that act against other target-force elements. Best results will be obtained in stiff, redundant structures.

Target-force loads can be applied at the same time as other loads. However, better convergence behavior may be obtained by applying target-force loads in a separate stage or Load Case when this is possible.

For more information:

- See Topic “Target-Force Load” (page 142) in Chapter “The Frame Element.”
- See Topic “Target-Force Load” (page 174) in Chapter “The Cable Element.”
- See Topic “Target-Force Load” (page 331) in Chapter “Load Patterns.”

Nonlinear Time-History Analysis

Time-history analysis is a step-by-step analysis of the dynamical response of a structure to a specified loading that may vary with time. The analysis may be linear or nonlinear. The Chapter describes concepts that apply only to nonlinear time-history analysis. You should first read Chapter “Linear Time-History Analysis” (page 397) which describes concepts that apply to all time-history analyses.

Advanced Topics

- Overview
- Nonlinearity
- Loading
- Initial Conditions
- Time Steps
- Nonlinear Modal Time-History Analysis (FNA)
- Nonlinear Direct-Integration Time-History Analysis

Overview

Time-history analysis is used to determine the dynamic response of a structure to arbitrary loading. The dynamic equilibrium equations to be solved are given by:

$$\mathbf{K} \mathbf{u}(t) + \mathbf{C} \dot{\mathbf{u}}(t) + \mathbf{M} \ddot{\mathbf{u}}(t) = \mathbf{r}(t)$$

where \mathbf{K} is the stiffness matrix; \mathbf{C} is the damping matrix; \mathbf{M} is the diagonal mass matrix; \mathbf{u} , $\dot{\mathbf{u}}$, and $\ddot{\mathbf{u}}$ are the displacements, velocities, and accelerations of the structure; and \mathbf{r} is the applied load. If the load includes ground acceleration, the displacements, velocities, and accelerations are relative to this ground motion.

Any number of time-history Load Cases can be defined. Each time-history case can differ in the load applied and in the type of analysis to be performed.

There are several options that determine the type of time-history analysis to be performed:

- Linear vs. Nonlinear.
- Modal vs. Direct-integration: These are two different solution methods, each with advantages and disadvantages. Under ideal circumstances, both methods should yield the same results to a given problem.
- Transient vs. Periodic: Transient analysis considers the applied load as a one-time event, with a beginning and end. Periodic analysis considers the load to repeat indefinitely, with all transient response damped out.

In a nonlinear analysis, the stiffness, damping, and load may all depend upon the displacements, velocities, and time. This requires an iterative solution to the equations of motion.

Before reading this Chapter on nonlinear analysis, you should first read Chapter “Linear Time-History Analysis” (page 397) which describes concepts that apply to all time-history analyses

Nonlinearity

The following types of nonlinearity are available in SAP2000:

- Material nonlinearity
 - Various type of nonlinear properties in Link/Support elements
 - Tension and/or compression limits in Frame elements

- Hinges in Frame elements
- Geometric nonlinearity in Cable elements
- Material nonlinearity in layered Shell elements
- Geometric nonlinearity
 - P-delta effects
 - Large displacement effects

For nonlinear direct-integration time-history analysis, all of the available nonlinearities may be considered.

For nonlinear modal time-history analysis, only the nonlinear behavior of the Link/Support elements is included. If the modes used for this analysis were computed using the stiffness from the end of a nonlinear load case, all other types of nonlinearities are locked into the state that existed at the end of that nonlinear load case.

Loading

The application of load for nonlinear time-history analysis is identical to that used for linear time-history analysis. Please see Topic “Loading” (page 398) in Chapter “Linear Time-History Analysis” for more information.

Initial Conditions

The initial conditions describe the state of the structure at the beginning of a time-history case. These include:

- Displacements and velocities
- Internal forces and stresses
- Internal state variables for nonlinear elements
- Energy values for the structure
- External loads

The accelerations are not considered initial conditions, but are computed from the equilibrium equation.

For nonlinear analyses, you may specify the initial conditions at the start of the analysis. You have two choices:

- Zero initial conditions: the structure has zero displacement and velocity, all elements are unstressed, and there is no history of nonlinear deformation.
- Continue from a previous nonlinear analysis: the displacements, velocities, stresses, loads, energies, and nonlinear state histories from the end of a previous analysis are carried forward.

There are some restrictions when continuing from a previous nonlinear case:

- Nonlinear static and nonlinear direct-integration time-history cases may be chained together in any combination, i.e., both types of analysis are compatible with each other.
- Nonlinear modal time-history (FNA) cases can only continue from other FNA cases that use modes from the same modal Load Case.

When continuing from a previous case, all applied loads specified for the present Load Case are incremental, i.e., they are added to the loads already acting at the end of the previous case.

When performing a nonlinear time-history analysis, such as for earthquake loading, it is often necessary to start from a nonlinear static state, such as due to gravity loading. For nonlinear direct-integration analysis, you can continue from a nonlinear static Load Case. But since FNA analyses can only continue from other FNA cases, special consideration must be given to how to model static loading using FNA. See Topic “Nonlinear Modal Time-History Analysis (FNA)” (page 149) for more information.

Note that, by contrast, linear time-history analyses *always* start from zero initial conditions.

Time Steps

The choice of output time steps is the same for linear and nonlinear time-history analysis. Please see Topic “Time Steps” (page 402) in Chapter “Linear Time-History Analysis” for more information.

The nonlinear analysis will internally solve the equations of motion at each output time step and at each load function time step, just as for linear analysis. In addition, you may specify a maximum substep size that is smaller than the output time step in order to reduce the amount of nonlinear iteration, and also to increase the accuracy of direct-integration analysis. The program may also choose smaller substeps sizes automatically when it detects slow convergence.

Nonlinear Modal Time-History Analysis (FNA)

The method of nonlinear time-history analysis used in SAP2000 is an extension of the Fast Nonlinear Analysis (FNA) method developed by Wilson (Ibrahimbegovic and Wilson, 1989; Wilson, 1993). The method is extremely efficient, particularly for structural systems which are primarily linear elastic but which have a limited number of predefined nonlinear elements. However, there is no limit on the number of nonlinear elements that can be considered, provided that adequate modes are obtained. This is best done using a sufficient number of Ritz vectors.

For the FNA method, all nonlinearity is restricted to the Link/Support elements. This includes hinges that are modeled as Links, as described in Topic “Analysis Modeling” (page 161) of Chapter “Hinge Properties”.

A short description of the method follows. The dynamic equilibrium equations of a linear elastic structure with predefined nonlinear Link/Support elements subjected to an arbitrary load can be written as:

$$\mathbf{K}_L \mathbf{u}(t) + \mathbf{C} \dot{\mathbf{u}}(t) + \mathbf{M} \ddot{\mathbf{u}}(t) = \mathbf{r}_N(t) + \mathbf{r}(t)$$

where \mathbf{K}_L is the stiffness matrix for the linear elastic elements (all elements except the Links/Supports); \mathbf{C} is the proportional damping matrix; \mathbf{M} is the diagonal mass matrix; \mathbf{r}_N is the vector of forces from the nonlinear degrees of freedom in the Link/Support elements; \mathbf{u} , $\dot{\mathbf{u}}$, and $\ddot{\mathbf{u}}$ are the relative displacements, velocities, and accelerations with respect to the ground; and \mathbf{r} is the vector of applied loads. See Topic “Loading” (page 398) in Chapter “Linear Time-History Analysis” for the definition of \mathbf{r} .

Initial Conditions

See Topic “Initial Conditions” (page 149) in this Chapter for a general discussion of initial conditions.

Because FNA analyses can only continue from other FNA analyses, special consideration must be given to how you can model static loads that may act on the structure prior to a dynamic analysis.

It is actually very simple to perform static analysis using FNA. The load is applied quasi-statically (very slowly) with high damping. To define a quasi-static FNA analysis:

- Define a ramp-type time-history function that increases linearly from zero to one over a length of time that is long (say ten times) compared to the first period

of the structure, and then holds constant for an equal length of time. Call this function “RAMPQS”

- Define a nonlinear modal time-history (FNA) case:
 - Call this case “HISTQS”
 - Start from zero or another FNA case
 - Apply the desired Load Pattern(s) using function “RAMPQS”
 - Use as few or as many time steps as you wish, but make sure the total time is at least twice the ramp-up time of function “RAMPQS”
 - Use high modal damping, say 0.99

You can use case “HISTQS” as the initial conditions for other FNA cases.

This approach is particularly useful for nonlinear analysis where the behavior of certain Link/Support elements, especially the Gap, Hook and Friction types, is strongly dependent on the *total* force or displacement acting on the elements.

Link/Support Effective Stiffness

For the purposes of analysis, a linear **effective stiffness** is defined for each degree of freedom of the nonlinear elements. The effective stiffness at nonlinear degrees of freedom is arbitrary, but generally varies between zero and the maximum nonlinear stiffness of that degree of freedom.

The equilibrium equation can then be rewritten as:

$$\mathbf{K} \mathbf{u}(t) - \mathbf{C} \dot{\mathbf{u}}(t) - \mathbf{M} \ddot{\mathbf{u}}(t) = \mathbf{r}(t) - [\mathbf{r}_N(t) - \mathbf{K}_N \mathbf{u}(t)]$$

where $\mathbf{K} = \mathbf{K}_L + \mathbf{K}_N$, with \mathbf{K}_L being the stiffness of all the linear elements and for the linear degrees of freedom of the Link/Support elements, and \mathbf{K}_N being the linear effective-stiffness matrix for all of the nonlinear degrees of freedom.

See Chapter “The Link/Support Element—Basic” (page 251) for more information.

Mode Superposition

Modal analysis is performed using the full stiffness matrix, \mathbf{K} , and the mass matrix, \mathbf{M} . It is *strongly recommended* that the Ritz-vector method be used to perform the modal analysis.

Using standard techniques, the equilibrium equation can be written in modal form as:

$$\mathbf{\Omega}^2 \mathbf{a}(t) + \mathbf{D} \dot{\mathbf{a}}(t) + \mathbf{I} \ddot{\mathbf{a}}(t) = \mathbf{q}(t) + \mathbf{q}_N(t)$$

where $\mathbf{\Omega}^2$ is the diagonal matrix of squared structural frequencies given by:

$$\mathbf{\Omega}^2 = \mathbf{K}^{-1} \mathbf{M}^{-1}$$

\mathbf{D} is the modal damping matrix which is assumed to be diagonal:

$$\mathbf{D} = \mathbf{C}^{-1} \mathbf{M}^{-1}$$

\mathbf{I} is the identity matrix which satisfies:

$$\mathbf{I} \mathbf{M} = \mathbf{M}$$

$\mathbf{q}(t)$ is the vector of modal applied loads:

$$\mathbf{q}(t) = \mathbf{r}^T(t) \mathbf{r}(t)$$

$\mathbf{q}_N(t)$ is the vector of modal forces from the nonlinear elements:

$$\mathbf{q}_N(t) = \mathbf{r}_N^T(t) [\mathbf{r}_N(t) \mathbf{K}_N \mathbf{u}(t)]$$

$\mathbf{a}(t)$ is the vector of modal displacement amplitudes such that:

$$\mathbf{u}(t) = \mathbf{a}(t) \mathbf{\Phi}$$

and $\mathbf{\Phi}$ is the matrix of mode shapes.

It should be noted that, unlike linear dynamic analysis, the above modal equations are coupled. In general the nonlinear forces, $\mathbf{q}_N(t)$, will couple the modes since they are functions of the modal displacements, $\mathbf{a}(t)$.

It is important to recognize that the solution to these modal equations is dependent on being able to adequately represent the nonlinear forces by the modal forces, $\mathbf{q}_N(t)$. This is not automatic, but requires the following special considerations:

- Mass and/or mass moments of inertia should be present at all nonlinear degrees of freedom.
- The Ritz-vector method should be used to determine the Modes, unless *all* possible structural Modes are found using eigenvector analysis

- The Ritz starting load vectors should include a nonlinear deformation load for each independent nonlinear degree of freedom
- A sufficient number of Ritz-vectors should be sought to capture the deformation in the nonlinear elements completely

For more information:

- See Topic “Ritz-Vector Analysis” (page 349) in Chapter “Modal Analysis”.
- See Chapter “The Link/Support Element—Basic” (page 251).

Modal Damping

As for linear modal time-history analysis, the damping in the structure is modeled using uncoupled modal damping. Each mode has a damping ratio, **damp**, which is measured as a fraction of critical damping and must satisfy:

$$0 \leq \text{damp} \leq 1$$

Modal damping has two different sources, which are described in the following. Damping from these sources is added together. The program automatically makes sure that the total is less than one.

Important note: For linear modal time-history analysis, the linear effective damping for the Link/Support elements is also used. However, it is not used for nonlinear modal time-history analysis.

Modal Damping from the Load Case

For each nonlinear modal time-history Load Case, you may specify modal damping ratios that are:

- Constant for all modes
- Linearly interpolated by period or frequency. You specify the damping ratio at a series of frequency or period points. Between specified points the damping is linearly interpolated. Outside the specified range, the damping ratio is constant at the value given for the closest specified point.
- Mass and stiffness proportional. This mimics the proportional damping used for direct-integration, except that the damping value is never allowed to exceed unity.

In addition, you may optionally specify damping overwrites. These are specific values of damping to be used for specific modes that replace the damping obtained by one of the methods above. The use of damping overwrites is rarely necessary.

It is also important to note that the assumption of modal damping is being made with respect to the total stiffness matrix, \mathbf{K} , which includes the effective stiffness from the nonlinear elements. If non-zero modal damping is to be used, then the effective stiffness specified for these elements is important. The effective stiffness should be selected such that the modes for which these damping values are specified are realistic.

In general it is recommended that either the initial stiffness of the element be used as the effective stiffness or the secant stiffness obtained from tests at the expected value of the maximum displacement be used. Initially-open gap and hook elements and all damper elements should generally be specified with zero effective stiffness.

Composite Modal Damping from the Materials

Modal damping ratios, if any, that have been specified for the Materials are converted automatically to composite modal damping. Any cross coupling between the modes is ignored. These modal-damping values will generally be different for each mode, depending upon how much deformation each mode causes in the elements composed of the different Materials.

Iterative Solution

The nonlinear modal equations are solved iteratively in each time step. The program assumes that the right-hand sides of the equations vary linearly during a time step, and uses exact, closed-form integration to solve these equations in each iteration. The iterations are carried out until the solution converges. If convergence cannot be achieved, the program divides the time step into smaller substeps and tries again.

Several parameters are available for you to control the iteration process. In general, the use of the default values is recommended since this will solve most problems. If convergence cannot be achieved, inaccurate results are obtained, or the solution takes too long, changing these control parameters may help. However, you should first check that reasonable loads and properties have been specified, and that appropriate Modes have been obtained, preferably using the Ritz vector method.

The parameters that are available to control iteration and substepping are:

- The relative force convergence tolerance, **ftol**

- The relative energy convergence tolerance, **etol**
- The maximum allowed substep size, **dtmax**
- The minimum allowed substep size, **dtmin**
- The maximum number of force iterations permitted for small substeps, **itmax**
- The maximum number of force iterations permitted for large substeps, **itmin**
- The convergence factor, **cf**

These parameters are used in the iteration and substepping algorithm as described in the following.

Force Convergence Check

Each time step of length **dt** is divided into substeps as needed to achieve convergence. In each substep, the solution is iterated until the change in the right-hand side of the modal equations, expressed as a fraction of the right-hand side, becomes less than the force tolerance, **ftol**. If this does not occur within the permitted number of iterations, the substep size is halved and the iteration is tried again.

The default value for **ftol** is 10^{-5} . It must satisfy **ftol** > 0.

Energy Convergence Check

If force convergence occurs within the permitted number of iterations, the work done by the nonlinear forces is compared with the work done by all the other force terms in the modal equilibrium equations. If the difference, expressed as a fraction of the total work done, is greater than the energy tolerance, **etol**, the substep size is halved and the iteration is tried again.

This energy check essentially measures how close to linear is the variation of the nonlinear force over the time step. It is particularly useful for detecting sudden changes in nonlinear behavior, such as the opening and closing of gaps or the onset of yielding and slipping. Setting **etol** greater than unity turns off this energy check.

The default value for **etol** is 10^{-5} . It must satisfy **etol** > 0.

Maximum and Minimum Substep Sizes

If the substep meets both the force and energy convergence criteria, the results of the substep are accepted, and the next substep is attempted using twice the previous substep length. The substep size is never increased beyond **dtmax**.

When the substep size is halved because of failure to meet either the force or energy convergence criteria, the resulting substep size will never be set less than **dtmin**. If the failed substep size is already **dtmin**, the results for the remaining time steps in the current History are set to zero and a warning message is issued.

The default value for **dtmax** is **dt**. The default value for **dtmin** is $\text{dtmax} \cdot 10^{-9}$. They must satisfy $0 < \text{dtmin} < \text{dtmax} < \text{dt}$.

Maximum Number of Iterations

The maximum number of iterations permitted for force iteration varies between **itmin** and **itmax**. The actual number permitted for a given substep is chosen automatically by the program to achieve a balance between iteration and substepping. The number of iterations permitted tends to be larger for smaller substeps.

The default values for **itmin** and **itmax** are 2 and 100, respectively. They must satisfy $2 < \text{itmin} < \text{itmax}$.

Convergence Factor

Under-relaxation of the force iteration may be used by setting the convergence factor, **cf**, to a value less than unity. Smaller values increase the stability of the iteration, but require more iterations to achieve convergence. This is generally only needed when Damper-type elements are present with nonlinear damping exponents. Specifying **cf** to be greater than unity may reduce the number of iterations required for certain types of problems, but may cause instability in the iteration and is not recommended.

The default value for **cf** is 1. It must satisfy $\text{cf} > 0$.

Static Period

Normally all modes are treated as being dynamic. In each time step, the response of a dynamic mode has two parts:

- Forced response, which is directly proportional to the modal load
- Transient response, which is oscillatory, and which depends on the displacements and velocities of the structure at the beginning of the time step

You may optionally specify that high-frequency (short period) modes be treated as static, so that they follow the load without any transient response. This is done by specifying a static period, **tstat**, such that all modes with periods less than **tstat** are

considered to be static modes. The default for **tstat** is zero, meaning that all modes are considered to be dynamic.

Although **tstat** can be used for any nonlinear time-history analysis, it is of most use for quasi-static analyses. If the default iteration parameters do not work for such an analysis, you may try using the following parameters as a starting point:

- **tstat** greater than the longest period of the structure
- **itmax** = **itmin** 1000
- **dtmax** = **dtmin** = **dt**
- **ftol** 10^{-6}
- **cf** = 0.1

This causes all modes to be treated as static, and uses iteration rather than substepping to find a solution. The choice of parameters to achieve convergence is very problem dependent, and you should experiment to find the best values to use for each different model.

Nonlinear Direct-Integration Time-History Analysis

Direct integration of the full equations of motion without the use of modal superposition is available in SAP2000. While modal superposition is usually more accurate and efficient, direct-integration does offer the following advantages:

- Full damping that couples the modes can be considered
- Impact and wave propagation problems that might excite a large number of modes may be more efficiently solved by direct integration
- All types of nonlinearity available in SAP2000 may be included in a nonlinear direct integration analysis.

Direct integration results are extremely sensitive to time-step size in a way that is not true for modal superposition. You should always run your direct-integration analyses with decreasing time-step sizes until the step size is small enough that results are no longer affected by it.

Time Integration Parameters

See Topic “Linear Direct-Integration Time-History Analysis” (page 149) for information about time-integration parameters. The same considerations apply here as for linear analysis.

If your nonlinear analysis is having trouble converging, you may want to use the HHT method with **alpha** = -1/3 to get an initial solution, then re-run the analysis with decreasing time step sizes and **alpha** values to get more accurate results.

Nonlinearity

All material nonlinearity that has been defined in the model will be considered in a nonlinear direct-integration time-history analysis.

You have a choice of the type of geometric nonlinearity to be considered:

- None
- P-delta effects
- Large displacement effects

If you are continuing from a previous nonlinear analysis, it is strongly recommended that you select the same geometric nonlinearity parameters for the current case as for the previous case. See Chapter “Geometric Nonlinearity” (page 409) for more information.

Initial Conditions

See Topic “Initial Conditions” (page 149) in this Chapter for a general discussion of initial conditions.

You may continue a nonlinear direct-integration time-history analysis from a nonlinear static analysis or another direct-integration time-history nonlinear analysis. It is strongly recommended that you select the same geometric nonlinearity parameters for the current case as for the previous case.

When multiple Mass Sources have been defined, you can specify which Mass Source to use for a nonlinear static or nonlinear direct-integration time-history load case. However, the default is to use the same Mass Source as the previous load case, or the default Mass Source if starting from zero initial conditions. See Topic “Mass Source” (page 334) in Chapter “Load Patterns” for more information.

Damping

In direct-integration time-history analysis, the damping in the structure is modeled using a full damping matrix. Unlike modal damping, this allows coupling between the modes to be considered.

Direct-integration damping has two different sources, which are described in the following. Damping from these sources is added together.

Important note: For linear direct-integration time-history analysis, the linear effective damping for the nonlinear Link/Support elements is also used. However, it is not used for nonlinear direct-integration time-history analysis.

Proportional Damping from the Load Case

For each direct-integration time-history Load Case, you may specify proportional damping coefficients that apply to the structure as a whole. The damping matrix is calculated as a linear combination of the stiffness matrix scaled by a coefficient, c_K , and the mass matrix scaled by a second coefficient, c_M .

You may specify these two coefficients directly, or they may be computed by specifying equivalent fractions of critical modal damping at two different periods or frequencies.

For each nonlinear element in the structure, the coefficient c_K multiplies the initial stiffness matrix, i.e., the stiffness of the element at zero initial conditions, regardless of the current nonlinear state of the element. The exception to this rule is that if the current nonlinear state has zero stiffness *and* zero force or stress (such as an open gap or a cracked concrete material), then zero damping is assumed. In the case where the initial stiffness is different in the negative and positive direction of loading, the larger stiffness is used.

For cable elements, the damping matrix is proportional to the stiffness matrix for an equivalent truss element along the current chord having the same axial stiffness (AE/L), where L is the undeformed length.

Mass proportional damping is linearly proportional to period. It is related to the motion of the structure, as if the structure is moving through a viscous fluid. Mass proportional damping may excessively damp out long period components.

Proportional Damping from the Materials

You may specify stiffness and mass proportional damping coefficients for individual materials. For example, you may want to use larger coefficients for soil materials than for steel or concrete. The same interpretation of these coefficients applies as described above for the Load Case damping. Be sure not to double-count the damping by including the same amount in both the Load Case and the materials.

Nonlinear Solution Control

The nonlinear equations are solved using event-to-event stepping and/or equilibrium iteration in each time step. This may require re-forming and re-solving the stiffness and damping matrices at each step or iteration.

Several parameters are available for you to control the stepping and iteration process. These are described in the following. We recommend that you start with the default values of these parameters, except that you may want to vary the maximum substep size for reasons of accuracy.

If you are having convergence difficulties, you can try varying the solution control parameters below. However, you should also consider that the model itself may need improvement. Look for instabilities due to inadequate support, buckling, and excessively large stiffnesses. If you have hinges or materials that lose strength, make sure that this behavior is really necessary and that the negative slopes are not unrealistically too steep.

Maximum Substep Size

The analysis will always stop at every output time step, and at every time step where one of the input time-history functions is defined. You may, in addition, set an upper limit on the step size used for integration. For example, suppose your output time step size was 0.005, and your input functions were also defined at 0.005 seconds. If you set the Maximum Substep Size to 0.001, the program will internally take five integration substeps for every saved output time step. The program may automatically use even smaller substeps if necessary to achieve convergence when iterating.

The accuracy of direct-integration methods is very sensitive to integration time step, especially for stiff (high-frequency) response. You should try decreasing the maximum substep size until you get consistent results. You can keep your output time step size fixed to prevent storing excessive amounts of data.

Minimum Substep Size

When using iteration and the solution cannot converge within the specified maximum number of iterations, the program automatically reduces the current step size and tries again. You can limit the smallest substep size the program will use. If the program tries to reduce the step size below this limit, it will stop the analysis and indicate that convergence had failed. The minimum substep size has no effect when using event-to-event stepping without iteration.

Event-to-Event Stepping Control

Event-to-event stepping is available as an option for nonlinear direct-integration time-history analysis. Using this strategy, an increment of load is applied in each time step until one of the nonlinear elements experiences an “event” that causes a significant change in the stiffness of the structure. The stiffness matrix is then reformed, and another load increment is applied to reach the next event. This continues in a series of event steps until the entire load for that time step has been applied. If no event occurs within a time step, the full load for that time step is applied all at once.

Events are implemented for the following nonlinear models: all frame hinges; layered shells with directional or coupled nonlinear behavior; and links with multi-linear plasticity, gap, hook, friction-pendulum, and triple-pendulum behavior. Events typically include first yield, unloading, strength loss, increasing stiffness, and other significant changes in behavior that depend on the type of element or material.

When event-to-event stepping is used, you may specify whether or not to perform iteration to reduce any equilibrium error to within a specified convergence tolerance (see next subtopic). If iteration is requested, this will be performed at the end of each time step, not at the end of each event.

If iteration is not requested, the relative equilibrium error will be reported in the analysis log file (.LOG) and the unbalanced load will be applied in the next time step. Although the structure may be slightly out of equilibrium at each time step, the unbalance does not accumulate over time. This method can be quite effective because the deviation from linearity is usually small between events. However, iteration for equilibrium is recommended for models having significant geometric nonlinearity, including models with cable elements.

When iteration is not performed, there can be no convergence failure. For this reason, event-to-event stepping without iteration may help with certain load cases that otherwise fail to complete, although the results should be reviewed carefully for acceptable equilibrium error.

You can specify whether or not to use event-to-event stepping. When events are used, each load increment tends to be more linear, thus reducing equilibrium unbalance and lessening the need for iteration. On the other hand, models with a large number of nonlinear elements and/or hinges may generate a large number of events, which can significantly increase the analysis time.

When using event stepping, you can specify a relative event-lumping tolerance that will group multiple events together. Effectively, the load increment calculated to reach the first event is increased by the tolerance to include other events that would

occur soon afterward. Larger event tolerances reduce the number of event steps taken, but may increase the equilibrium unbalance or amount of iteration required.

Note that for isotropic frame hinges the event lumping tolerance may also allow the force or deformation response to deviate from the backbone curve. The relative magnitude of the deviations is on the order of the relative lumping tolerance.

You can also specify the maximum number of events to allow per time step. Increasing this number may increase the accuracy of the analysis when iteration is not used, or lessen the need for iteration when used. On the other hand, allowing fewer events per time step may increase the speed of the analysis at the expense of more equilibrium error or iteration.

Whether or not to use event stepping is a problem-dependent decision. Normally using events will improve convergence behavior. However, if a large number of event steps are generated that make the analysis run slowly, try turning event-stepping off for that model or load case. Either way, results for a stable model should be consistent to within the convergence tolerance when iteration is used.

Nonlinear Iteration

Iteration may be used to ensure that equilibrium is achieved at each step of the analysis to within a specified convergence tolerance. You may specify whether or not to use iteration when using event-to-event stepping, as described in the preceding subtopic. Iteration is always used when event stepping is not used.

When using iteration, you can specify the relative convergence tolerance used to compare the magnitude of force error with the magnitude of the force acting on the structure. You may need to use significantly smaller values of convergence tolerance to get good results for large-displacements problems than for other types of nonlinearity. Try decreasing this value until you get consistent results.

For each time step, constant-stiffness iteration is tried first. If convergence is not achieved, Newton-Raphson (tangent-stiffness) iteration is tried next. If both fail, the step size is reduced, and the process is repeated.

You can separately control the number of constant-stiffness and Newton-Raphson iterations allowed in each step. Setting either parameter to zero prevents that type of iteration. Setting both to zero causes the analysis to use the default number of iterations. Constant-stiffness iterations are faster than Newton-Raphson iterations, but the latter are usually more effective, especially for cables and geometric nonlinearity. The default values work well in many situations.

Line Search Option

Line search is an iteration option that can increase the efficiency of equilibrium iterations by scaling the solution increment in a trial-and-error fashion to find the smallest unbalance. This increases the computation time for each iteration, but often results in fewer iterations and better convergence behavior, with a net gain in efficiency. Line search is particularly effective for stiffening systems, such as tightening cables, closing gaps, unloading from plastic yielding, and sticking after frictional sliding.

You can choose whether or not to use line search, provided that iteration is being used. Unlike nonlinear static analysis, line search for time-history analysis can be used in conjunction with event-to-event stepping.

When using line search, specify the following parameters:

- **Maximum number of line searches per Iteration.** Use this parameter to limit the number of line searches allowed per iteration, usually in the range of 5 to 40.
- **Line-Search Acceptance Tolerance (Relative).** The solution increment is successively reduced or increased until the minimum error is found, the maximum number of line searches is reached, or the relative change in error from the previous trial is less than the acceptance tolerance set here. The practical range is about 0.02 to 0.50. This value should not be too small, since the goal is only to improve the next iteration.
- **Line-Search Step Factor.** For each line-search trial, the solution increment is successively reduced by the step factor specified here until the minimum error is found, the maximum number of line searches is reached, or the acceptance tolerance is satisfied. If reducing the increment does not decrease the error, the increment is instead increased by the step factor following the same procedure. This factor must be greater than 1.0, and should usually be no more than 2.0.

The default values are recommended as starting point. For stiffening systems that are slow to converge, try increasing the maximum number of line searches per iteration, decreasing the acceptance tolerance, and reducing the step factor.

Frequency-Domain Analyses

Frequency-domain analysis is based upon the dynamical response of the structure to harmonically varying load. Two types of frequency-domain Load Cases are currently available: deterministic Steady-State analysis and probabilistic Power-Spectral-Density analysis.

Advanced Topics

- Overview
- Harmonic Motion
- Frequency Domain
- Damping
- Loading
- Frequency Steps
- Steady-State Analysis
- Power-Spectral-Density Analysis

Overview

Frequency-domain analysis is based upon the dynamical response of the structure to harmonically varying load. The analysis is performed at one or more frequencies of vibration. At each frequency, the loading varies with time as sine and cosine functions. Two types of frequency-domain Load Cases are currently available: steady-state analysis and power-spectral-density analysis.

Steady-state analysis computes the deterministic response at each requested frequency. The loading may have components at acting different phase angles. The phase angles of the response are computed and may be displayed.

Power-spectral-density analysis is based on a probabilistic spectrum of loading. The analysis computes a probabilistic spectrum for each response component. In addition, a single probabilistic expected value for each response component is presented that is the root-mean-square (RMS) of the probabilistic spectrum. The loading may have components at acting different phase angles, but the phase information is not preserved for the probabilistic response.

Hysteretic and viscous damping may be specified. Frequency-dependent properties for Link and Support elements, if defined, are considered in the analyses. All analyses are performed in the complex domain.

Harmonic Motion

Harmonic loading is of the form $r(t) = p\theta \cos(\omega t) + p90 \sin(\omega t)$, where ω is the circular frequency of the excitation. This loading is assumed to exist for all time, so that transient components of the response have vanished. In other words, steady-state conditions have been achieved.

The spatial loading consists of two parts: the in-phase component $p\theta$, and the 90 out-of-phase component $p90$. The spatial distributions do not vary as a function of time.

The equilibrium equations for the structural system are of the following form:

$$K u(t) + C \dot{u}(t) + M \ddot{u}(t) = r(t) = p\theta \cos(\omega t) + p90 \sin(\omega t)$$

where K is the stiffness matrix, C is the viscous damping matrix, M is the diagonal mass matrix, and u , \dot{u} , and \ddot{u} are the joint displacements, velocities, and accelerations, respectively.

Frequency Domain

It is more convenient to re-write the equations in complex form. The loading is then given by

$$\bar{r}(t) = \bar{p} \exp(i \omega t) = \bar{p} (\cos(\omega t) + i \sin(\omega t))$$

where the overbar indicates a complex quantity. The real cosine term represents the in-phase component, and the imaginary sine term represents the 90° out-of-phase component.

The steady-state solution of this equation requires that the joint displacements be of the same form:

$$\bar{u}(t) = \bar{a} \exp(i \omega t) = \bar{a} (\cos(\omega t) + i \sin(\omega t))$$

Substituting these into the equation of motion yields:

$$[\mathbf{K} - i \omega \mathbf{C} - \omega^2 \mathbf{M}] \bar{a} = \bar{p}$$

We can define a complex **impedance** matrix

$$\bar{\mathbf{K}} = \mathbf{K} - i \omega \mathbf{C} - \omega^2 \mathbf{M}$$

where the real part represents stiffness and inertial effects, and the imaginary part represents damping effects. Note that the real part may be zero or negative. The equations of motion can be written:

$$\bar{\mathbf{K}}(\omega) \bar{a}(\omega) = \bar{p}(\omega) \quad (\text{Eqn. 1})$$

where here we emphasize that the impedance matrix, the loading, and the displacements are all functions of frequency.

The impedance matrix is a function of frequency not only because of the inertial and damping terms, but also because frequency-dependent Link and Support element properties are permitted. Thus

$$\bar{\mathbf{K}}(\omega) = \mathbf{K}(\omega) - i \omega \mathbf{C}(\omega) - \omega^2 \mathbf{M}(\omega)$$

Please see Topic “Frequency-Dependent Link/Support Properties” (page 306) in Chapter “The Link/Support Element—Advanced” for more information.

Damping

It is common for frequency-domain problems to specify a hysteretic (displacement-based) damping matrix D rather than a viscous (velocity-based) damping matrix C . These are related as:

$$D = C$$

[As an aside, note that from this definition a nonzero value of hysteretic damping D at $\omega = 0$ (static conditions) results in an undefined value for viscous damping C . This leads to considerations of the noncausality of hysteretic damping, as discussed in Makris and Zhang (2000). However, this is usually ignored.]

Hysteretic damping may be specified as a function of frequency, i.e., $D = D(\omega)$, and there is no restriction imposed on the value at $\omega = 0$.

Using hysteretic damping, the complex impedance matrix becomes

$$\bar{K}(\omega) = K(\omega) - \omega^2 M - iD(\omega)$$

Sources of Damping

In frequency-domain analysis, the damping matrix $D(\omega)$ has four different sources as described in the following. Damping from these sources is added together.

Hysteretic Damping from the Load Case

For each Steady-state or Power-spectral-density Load Case, you may specify proportional damping coefficients that apply to the structure as a whole. The damping matrix is calculated as a linear combination of the stiffness matrix scaled by a coefficient, d_K , and the mass matrix scaled by a second coefficient, d_M . For most practical cases, $d_M = 0$ and only d_K is used.

You may specify that these coefficients are constant for all frequencies, or they may be linearly interpolated between values that you specify at different frequencies. Thus the damping matrix becomes:

$$D(\omega) = d_K(\omega)K + d_M(\omega)M$$

The stiffness matrix used here includes all elements in the structure except for any Link or Support elements that have frequency-dependent properties.

If you specify $d_M = 0$ and $d_K(\omega)$ to linearly increase with frequency, this is equivalent to viscous damping.

You can approximate modal damping by setting $d_M = 0$ and $d_K(\omega) = 2d(\omega)$, where $d(\omega)$ is the modal damping ratio. For example, if you typically use a constant 5% modal damping for all modes, the equivalent hysteretic damping value is a constant $d_K(\omega) = 0.10$. For each mode, this leads to approximately the same level of response at resonance.

Hysteretic Damping from the Materials

You may specify stiffness and mass proportional damping coefficients for individual materials. For example, you may want to use larger coefficients for soil materials than for steel or concrete. The same interpretation of these coefficients applies as described above for the Load Case damping. Be sure not to double-count the damping by including the same amount in both the Load Case and the materials.

Hysteretic Damping from Frequency-Dependent Link/Support Elements

For any Link or Support elements that have frequency-dependent properties assigned to them, the frequency-dependent hysteretic damping values are added to the damping matrix used.

Viscous Damping from the Link/Support Elements

For any Link or Support elements that *do not* have frequency-dependent properties assigned to them, the linear effective-damping coefficients are multiplied by frequency ω and added to the damping matrix. The linear effective-damping values are used regardless of whether or not nonlinear damping coefficients have been specified.

Loading

The load, $\bar{p}(\omega)$, applied in a given Steady-state or Power-spectral-density case may be an arbitrary function of space and frequency. It can be written as a finite sum of spatial load vectors, p_j , multiplied by frequency functions, $f_j(\omega)$, as:

$$\bar{p}(\omega) = \sum_j s_j f_j(\omega) p_j e^{i\phi_j} = \sum_j s_j f_j(\omega) p_j (\cos \phi_j + i \sin \phi_j) \quad (\text{Eqn. 2})$$

Here s_j is a scale factor and ϕ_j is the phase angle at which load p_j is applied.

The program uses Load Patterns and/or Acceleration Loads to represent the spatial load vectors, p_j , as described below.

The frequency functions used here depend on the type of analysis. See Topics “Steady-State Analysis” (page 471) and “Power-Spectral-Density Analysis” (page 473) in this Chapter for more information.

If Acceleration Loads are used, the displacements, velocities, and accelerations are all measured relative to the ground. The frequency functions associated with the Acceleration Loads m_x , m_y , and m_z are the corresponding components of uniform ground acceleration, \ddot{u}_{gx} , \ddot{u}_{gy} , and \ddot{u}_{gz} .

See Topic “Functions” (page 350) in Chapter “Load Cases” for more information.

Defining the Spatial Load Vectors

To define the spatial load vector, p_j , for a single term of the loading sum of Equation 2, you may specify either:

- The label of a Load Pattern using the parameter **load**, or
- An Acceleration Load using the parameters **csys**, **ang**, and **acc**, where:
 - **csys** is a fixed coordinate system (the default is zero, indicating the global coordinate system)
 - **ang** is a coordinate angle (the default is zero)
 - **acc** is the Acceleration Load (U1, U2, or U3) in the acceleration local coordinate system as defined below

Each Acceleration Load in the loading sum may have its own acceleration local coordinate system with local axes denoted 1, 2, and 3. The local 3 axis is always the same as the Z axis of coordinate system **csys**. The local 1 and 2 axes coincide with the X and Y axes of **csys** if angle **ang** is zero. Otherwise, **ang** is the angle from the X axis to the local 1 axis, measured counterclockwise when the +Z axis is pointing toward you. This is illustrated in Figure 82 (page 400).

The acceleration local axes are always referred to as 1, 2, and 3. The global Acceleration Loads m_x , m_y , and m_z are transformed to the local coordinate system for loading.

It is generally recommended, but not required, that the same coordinate system be used for all Acceleration Loads applied in a given Load Case.

Load Patterns and Acceleration Loads may be mixed in the loading sum.

For more information:

- See Chapter “Load Patterns” (page 321).

- See Topic “Acceleration Loads” (page 338) in Chapter “Load Patterns”.

Frequency Steps

Frequency-domain analyses are performed at discrete frequency steps. For a Steady-state or Power-spectral-density Load Case, you may request the response at the following frequencies:

- A required range of equally spaced frequencies. This is defined by specifying the first frequency, $f_1 = 0$; the last frequency, $f_2 = f_1$; and the number of increments, $n = 0$. This results in the following set of frequencies:

$$f_1, f_1 + f, f_1 + 2f, \dots, f_2,$$

where $f = (f_2 - f_1) / n$.

- Optionally, at all frequencies calculated in a specified Modal Load Case. Only frequencies that fall within the frequency range f_1 to f_2 will be used. See Chapter “Modal Analysis” (page 365) for more information.
- Optionally, at specified fractional deviations from all frequencies calculated in a specified Modal Load Case. For example, suppose you specify fractional deviations of 0.01 and -0.02. For each frequency f found by the Modal Load Case, the frequency-domain analysis will be performed at $1.01f$ and $0.98f$. Only frequencies that fall within the frequency range f_1 to f_2 will be used.
- Optionally, at any number of directly specified frequencies f . Only frequencies that fall within the frequency range f_1 to f_2 will be used.

Frequencies may be specified in Hz (cycles/second) or RPM (cycles/minute). These will be converted to circular frequencies, ω , by the program.

The use of modal frequencies and their fractional deviations can be very important to capture resonant behavior in the structure. Any set of equally-spaced frequencies could easily skip over the most significant response in a given frequency range. The use of directly specified frequencies can be important when you are concerned about sensitive equipment that may respond strongly at certain frequencies.

Steady-State Analysis

Steady-state analysis seeks the response of the structure at one or more frequencies to loading of the form:

$$\bar{p}(\omega) = \sum_j s_j f_j(\omega) \mathbf{p}_j e^{i\phi_j} = \sum_j s_j f_j(\omega) \mathbf{p}_j (\cos \phi_j + i \sin \phi_j)$$

See Topic “Loading” (page 469) in this Chapter for more information about this type of loading.

The frequency function, $f_j(\omega)$, is given directly by a steady-state function that you define. It represents the magnitude of the load before scaling by s_j . See Topic “Functions” (page 350) in Chapter “Load Cases” for more information.

If you are interested in the response characteristics of the structure itself, a constant function could be used, i.e., $f_j(\omega) = 1$. For loading that is caused by rotating machinery, $f_j(\omega) = \omega^2$ might be used.

Example

Suppose we have a machine with a spinning flywheel that has an eccentric mass. The mass is m and the center of mass is eccentric by an amount e . The flywheel spins about an axis parallel to the global Y axis. This machine is mounted on a structure, and we are interested in the steady-state response of the structure to the machine running at any speed in the range from 0 to 30Hz (1800 RPM).

The magnitude of the force from the eccentric mass that acts on the center of rotation is given by $em\omega^2$. This force rotates in the X-Z plane. To define the loading, we need the following:

- A Load Pattern, say “ECCX”, in which a unit load in the +X direction is assigned to the joint that represents the center of the flywheel.
- Another Load Pattern, say “ECCZ”, in which a unit load in the +Z direction is assigned to the same joint.
- A Steady-state Function, say “FSQUARED”, which varies as $f_j(\omega) = \omega^2$
- A Modal Load Case, say “MODAL”, which calculates all natural frequencies in the range from 0 to 30Hz. This can be for eigen or Ritz vectors; if Ritz, use the two Load Patterns “ECCX” and “ECCZ” as the starting load vectors.

We then define a Steady-state Load Case with the following features:

- The loads applied include the following two contributions:
 - Load Pattern “ECCX”, Function “FSQUARED”, a scale factor equal to em , and a phase angle of 0

- Load Pattern “ECCZ”, Function “FSQUARED”, a scale factor equal to em , and a phase angle of 90°
- Frequency range from 0 to 30Hz, with 15 increments (every 2Hz)
- Additional frequencies at the modal frequencies of Load Case “MODAL”
- Additional frequencies at the following fractional deviations from the modal frequencies of Load Case “MODAL”: +0.01, 0.01, +0.02, 0.02, +0.03, 0.03, +0.05, 0.05
- Stiffness-proportional hysteretic damping with a coefficient of $d_k = 0.04$, corresponding to modal damping of 2%, which may be appropriate for small vibrations

After analysis, we can plot the deflected shape or force/stress response at any of the requested frequencies and at any phase angle. For example, the response at phase angle 0° primarily represents the behavior due to horizontal loading, plus a damping component due to vertical loading. We can instead plot the magnitude of the response at any requested frequency, where the magnitude is given by the square-root of the sum of the squares of the real (0°) and imaginary (90°) response components.

It is also possible to display plots of any response quantity as a function of frequency, yielding a frequency spectrum. This can be done for the component at any phase angle, or for the magnitude of the response.

Power-Spectral-Density Analysis

Power-spectral-density (PSD) analysis is similar to Steady-state analysis in that it considers the harmonic behavior of the structure over a range of frequencies. However, the loading is considered to be probabilistic over the frequency range of the analysis, and so too is the response. This probabilistic response can be integrated over the frequency range to determine a single expected value. This can be useful, for example, for fatigue design.

A PSD Load Case considers *correlated* loading of the form:

$$\bar{p}(t) = \sum_j s_j f_j(t) p_j e^{i \omega_j t} + \sum_j s_j f_j(t) p_j (\cos \omega_j t + i \sin \omega_j t)$$

See Topic “Loading” (page 469) in this Chapter for more information about this type of loading.

The frequency function, $f_j(\omega)$, used in this sum is taken as the *square root* of a Power-spectral-density function that you define.

To explain this further, PSD functions are specified as load-squared per unit of frequency. In order to combine correlated loading algebraically, the square-root of these functions are used. Normally one would expect that the same PSD function would be used for all correlated loading terms, but this is not required. Note that using a scale factor $s_j = 2$ in the sum here is the same as multiplying the PSD function itself by a factor of four. See Topic “Functions” (page 350) in Chapter “Load Cases” for more information.

The PSD curve for any response quantity (displacement, force, stress, etc.) is given by the square of the magnitude of that calculated response, plotted at every requested frequency step. The square-root of the integral under the PSD curve for a given response quantity gives the probabilistic expected value for that quantity, i.e., the root-mean-square (RMS) value. This will always be a positive number.

Because the PSD curves represent the square of the response, most of the integrated area will be near resonant frequencies of the structure. For accuracy, it is very important to capture the response at frequency steps at and around the natural modes of the structure.

Uncorrelated loading should be defined in separate PSD Load Cases, and then combined using SRSS-types of Load Combinations. See Topic “Load Combinations (Combos)” (page 351) in Chapter “Load Cases.”

Example

Consider the same example used in Topic “Steady-State Analysis” (page 471) of this Chapter. Suppose that the machine is expected to operate 95% of the time in the range of 20 to 25Hz, and 5% of the time at other frequencies from 0 to 30Hz.

The only difference between the definition of the two types of Load Cases for this problem is in the functions. Now we will use a PSD Function, say “FPOWER”, defined as follows:

$$F_j(\omega) = \begin{cases} 0.05 / 25\text{Hz}, & 0 < \omega < 20\text{Hz} \\ 0.95 / 5\text{Hz}, & 20 < \omega < 25\text{Hz} \\ 0.05 / 25\text{Hz}, & 25 < \omega < 30\text{Hz} \end{cases}$$

Note that the 2 term is squared again. However, in the definition of the PSD Load Case, the scale factor will still be em (not squared), since it was not included in the PSD function itself.

The two loads, “ECCX” and “ECCZ”, must be combined in the same Load Case because they are clearly correlated. However, if a second machine with its own independent functioning was mounted to the same structure, this should be analyzed in a separate PSD Load Case and the two cases combined in an SRSS Load Combination.

Moving-Load Analysis

Moving-load analysis is available in CSiBridge to compute influence lines and surfaces for traffic lanes on bridge structures and to analyze these structures for the response due to vehicle live loads. A simplified version of moving-load analysis is available in SAP2000 for moving vehicles, such as cranes, on rails or other one-dimensional paths.

Advanced Topics

- Overview for CSiBridge
- Moving-Load Analysis in SAP2000
- Bridge Modeler
- Bridge Analysis Procedure
- Lanes
- Influence Lines and Surfaces
- Vehicle Live Loads
- General Vehicle
- Vehicle Response Components
- Standard Vehicles
- Vehicle Classes

- Moving Load Load Cases
- Moving Load Response Control
- Step-by-Step Analysis
- Computational Considerations

Overview for CSiBridge

For CSiBridge, moving-load analysis can be used to determine the response of bridge structures due to the weight of Vehicle live loads. Considerable power and flexibility is provided for determining the maximum and minimum displacements, forces, and stresses due to multiple-lane loads on complex structures, such as highway interchanges. The effects of Vehicle live loads can be combined with static and dynamic loads, and envelopes of the response can be computed.

The bridge to be analyzed can be created using the parametric Bridge Modeler in CSiBridge; built manually using Frame, Shell, Solid, and/or Link elements; or by combining these two approaches. The superstructure can be represented by a simple “spine” (or “spline”) model using Frame elements, or it can be modeled in full 3-dimensional detail using Shell or Solid elements.

Lanes are defined that represent where the live loads can act on the superstructure. Lanes may have width and can follow any straight or curved path. Multiple Lanes need not be parallel nor of the same length, so that complex traffic patterns may be considered. The program automatically determines how the Lanes load the superstructure, even if they are eccentric to a spine model. Conventional influence lines and surfaces due to the loading of each Lane can be displayed for any response quantity.

You may select Vehicle live loads from a set of standard highway and railway Vehicles, or you may create your own Vehicle live loads. Vehicles are grouped in Vehicle Classes, such that the most severe loading of each Class governs.

Each vehicle may generate vertical and/or horizontal forces, representing the effects of gravity, braking and acceleration, and centrifugal forces. Overturning moments due to super-elevation and centrifugal load may also be included.

Two types of live-load analysis can be performed:

- Influence-based enveloping analysis: Vehicles move in both directions along each Lane of the bridge. Using the influence surface, Vehicles are automatically located at such positions along the length and width of the Lanes to pro-

duce the maximum and minimum response quantities throughout the structure. Each Vehicle may be allowed to act on every lane or be restricted to certain lanes. The program can automatically find the maximum and minimum response quantities throughout the structure due to placement of different Vehicles in different Lanes. For each maximum or minimum extreme response quantity, the corresponding values for the other components of response can also be computed. Vertical and/or horizontal loading may be considered in any combination.

- Step-by-step analysis: Any number of Vehicles can be run simultaneously on the Lanes, each with its own starting time, position, direction and speed. Step-by-step static or time-history analysis can be performed, with nonlinear effects included if desired. Currently only vertical loading may be considered in step-by-step analysis.

For most design purposes the enveloping-type analysis using Moving-Load Load Cases is most appropriate. For special studies and unusual permit vehicles, the step-by-step approach can be valuable.

The moving-load features in SAP2000 are a subset of those available in CSiBridge, as described in the next topic.

Moving-Load Analysis in SAP2000

This chapter is primarily written to describe the full moving-load capabilities for bridge design as available in CSiBridge. Moving-load analysis is also available in SAP2000, but is intended for a different class of problems, primarily for cranes riding on rails and similar one-dimensional applications.

For SAP2000, the following differences apply:

- Lanes are called “Paths” in SAP2000.
- Paths are defined by reference to a sequence (line or curve) of Frame elements, and the vehicle loads will be applied to these elements.
- Paths may be defined that are eccentric to the specified frame elements.
- Width effects for Paths and Vehicles are not included.
- Influence lines, but not surfaces, are available.
- Only the general vehicle is available, with the following features:
 - Any number of concentrated axle and uniform inter-axle loads
 - One inter-axle distance may be variable

- No floating axle loads
- No loaded-length effects
- Vertical loading only
- No limitations on the type of response calculated for the vehicle
- The vehicle may be restricted to the length of the lane.

Bridge Modeler

The Bridge Modeler in CSiBridge provides a powerful way to create and manage simple or complex bridge models. The bridge is represented parametrically with a set of high-level objects: layout (alignment) lines, bents (pier supports), abutments (end supports), deck cross sections, prestress tendons, etc.

These objects are combined into a super object called a Bridge Object. Typically a single Bridge Object represents the entire structure, although you may need multiple Bridge Objects if you have parallel structures, or want to consider merges or splits.

A Bridge Wizard is available within the Bridge Modeler to guide you through the process of creating a bridge model, and help is available within the wizard itself.

The important thing to understand here is that this parametric model of the bridge exists independently from the discretization of the model into elements. You may choose to discretize the Bridge Object as Frames (spine model), Shells, or Solids, and you may choose the size of the elements to be used. You can change the discretization at any time without affecting your parameterized bridge model. When you do this, the previously generated elements are automatically deleted, and new elements created.

You can add additional elements to the model to represent features of the bridge that may not be provided by the Bridge Modeler. These elements will not be affected by changes to the Bridge Object or its discretization, although it may be necessary to move or modify them if you change the geometry of the bridge.

You can make changes to the elements generated from a Bridge Object, such as assigning different properties or additional loads. These changes will survive regeneration of the model if a new element is generated in exactly the same location. However, this may not occur if there are changes in bridge geometry or discretization, so it is best to check a regenerated model and make your changes again if necessary.

Moving-Load Analysis Procedure

There are two types of vehicle live-load analysis that can be performed: influence-base enveloping analysis, and step-by-step analysis with full correspondence. The basic steps required for these two types of analysis are as follows.

For both types of analysis:

- (1) Create a structural model using the Bridge Modeler and/or standard structural model-building techniques.
- (2) Define Lanes that specify the location where vehicles can travel.
- (3) Define Vehicles that represent the live load acting in the Lanes.

For Influence-Based Analysis:

- (4) Define Vehicle Classes that group together one or more Vehicles that should be enveloped.
- (5) Define Moving-Load Load Cases that specify which Vehicle Classes should be moved on which Lanes to produce the enveloped response, and what combination of vertical and horizontal loads are to be considered.
- (6) Specify Bridge Response parameters that determine for which elements moving-load response should be calculated, and set other parameters that control the influence-based analysis.
- (7) After running the analysis, you may view influence lines for any element response quantity in the structure, and envelopes of response for those elements requested under Bridge Response.

For Step-by-Step Analysis:

- (8) Define Load Patterns of type “Vehicle Live” that specify which Vehicles move on which Lanes, at what speed, and from what starting positions.
- (9) Apply the Vehicle-Live Load Patterns in Multi-Step Static Load Cases, or in Time-History Load Cases if you are interested in dynamical effects.
- (10) After running the analysis, you may view step-by-step response or envelopes of response for any element in the structure. You may create a video showing the step-by-step static or dynamic results. Influence lines are not available.

Both types of bridge analysis may exist in the same model. You may create additional Load Patterns and Load Cases, and combine the results of these with the results for either type of bridge analysis.

Lanes

The Vehicle live loads are considered to act in traffic Lanes transversely spaced across the bridge roadway. The number of Lanes and their transverse spacing can be chosen to satisfy the appropriate design-code requirements. For simple bridges with a single roadway, the Lanes will usually be parallel and evenly spaced, and will run the full length of the bridge structure.

For complex structures, such as interchanges, multiple roadways may be considered; these roadways can merge and split. Lanes need not be parallel nor be of the same length. The number of Lanes across the roadway may vary along the length to accommodate merges. Multiple patterns of Lanes on the same roadway may be created to examine the effect of different lateral placement of the Vehicles.

Lanes in SAP2000 are called Paths, and they have no width effects.

Centerline and Direction

A traffic Lane is defined with respect to a **reference line**, which can be either a bridge layout line or a line (path) of Frame elements. The transverse position of the Lane centerline is specified by its eccentricity relative to the reference line. Lanes are said to “run” in a particular direction, namely from the first location on the reference line used to define the Lane to the last.

Eccentricity

Each Lane across the roadway width will usually refer to the same reference line, but will typically have a different eccentricity. The eccentricity for a given Lane may also vary along the length.

The sign of a Lane eccentricity is defined as follows: in an elevation view of the bridge where the Lane runs from left to right, Lanes located in front of the roadway elements have positive eccentricity. Alternatively, to a driver traveling on the roadway in the direction that the Lane runs, a Lane to the right of the reference line has a positive eccentricity. The best way to check eccentricities is to view them graphically in the graphical user interface.

In a spine model, the use of eccentricities is primarily important for the determination of torsion in the bridge deck and transverse bending in the substructure. In Shell and Solid models of the superstructure, the eccentricity determines where the load is applied on the deck.

Centrifugal Radius

The radius is used to compute centrifugal loads, and is only available for lanes defined from a bridge layout line, not for those defined from a line (path) of frame elements. The centrifugal radius is specified independently of the actual geometry of the lane in order to give more control over the loading.

A positive radius curves to the right when traveling in the direction that the Lane runs, generating centrifugal forces that act to the left. Conversely, a negative radius curves to the left, generating forces that act to the right. A zero radius actually indicates infinite radius, or zero curvature, representing a straight Lane that generates no centrifugal forces. The radius may be constant or variable along the length of the Lane. When variable, the curvature (inverse of the radius) is interpolated linearly between specified values.

For the Eurocode and similar prescriptive codes, you should set any radii that are larger than the prescribed upper limit to zero so as to generate no centrifugal forces, and set radii that are smaller than the prescribed lower limit to the value of that lower limit.

Width

You may specify a width for each Lane, which may be constant or variable along the length of the Lane. When a Lane is wider than a Vehicle, each axle or distributed load of the Vehicle will be moved transversely in the Lane to maximum effect. If the Lane is narrower than the Vehicle, the Vehicle is centered on the Lane and the Vehicle width is reduced to the width of the Lane.

Interior and Exterior Edges

Certain AASHTO vehicles require that the wheel loads maintain a specified minimum distance from the edge of the lane. This distance may be different depending on whether the edge of the lane is at the edge of the roadway or is interior to the roadway. For each lane, you may specify for the left and right edges whether they are interior or exterior, with interior being the default. This only affects vehicles which specify minimum distances for the wheel loads. By default, vehicle loads

may be placed transversely anywhere in the lane, i.e., the minimum distance is zero. Left and right edges are as they would be viewed by a driver traveling in the direction the lane runs.

Discretization

A influence surface will be constructed for each Lane for the purpose of placing the vehicles to maximum effect. This surface is interpolated from unit point loads, called influence loads, placed along the width and length of the Lane. Using more influence loads increases the accuracy of the analysis at the expense of more computational time, memory, and disk storage.

You can control the number of influence loads by independently specifying the discretization to be used along the length and across the width of each Lane. Discretization is given as the maximum distance allowed between load points. Transversely, it is usually sufficient to use half the lane width, resulting in load points at the left, right, and center of the Lane. Along the length of the Lane, using eight to sixteen points per span is often adequate.

For Lanes applied to a general structure, the distance between the load points along the length is chosen as the largest value that provides uniform spacing but does not exceed the specified discretization length.

For Lanes applied to a bridge object created by the Bridge Modeler, the load points are created as follows:

- Along the Lane centerline, load points are placed at each discretization line. Discretization lines include the beginning and end of each span, in-span hinges, full-width diaphragms, user-specified span points, and other discretization lines that created so as not to exceed the bridge-object discretization length or angle.
- Between these load points, additional points are added at equal spaces that do not exceed the *minimum* value of:
 - The discretization length specified along-the-length for the Lane
 - The fraction of the lane length specified for the Lane, if any
 - The fraction of the span length specified for the Lane, if any, using the shortest span in the bridge object
 - The submesh size specified when updating the bridge object.

- For each load point located on the lane centerline, transverse load points are placed along the discretization line at equal distances not to exceed the discretization length specified across-the-width for the Lane.

As with analyses of any type, it is *strongly* recommended that you start with models that run quickly, using coarser discretization, so that you can gain experience with your model and perform reality checks. Later, you can increase the refinement until you achieve the desired level of accuracy and obtain the detailed results that you need.

Influence Lines and Surfaces

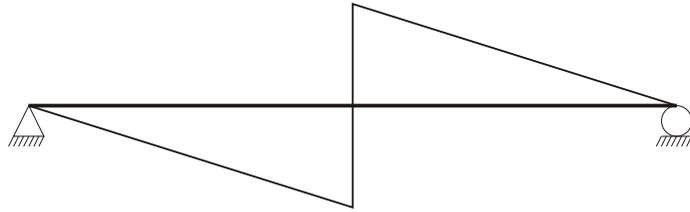
CSiBridge uses influence lines and surfaces to compute the response to vehicle live loads. Influence lines and surfaces are also of interest in their own right for understanding the sensitivity of various response quantities to traffic loads.

Influence lines are computed for Lanes of zero width, while influence surfaces are computed for Lanes having finite width. Only influence lines are available in SAP2000.

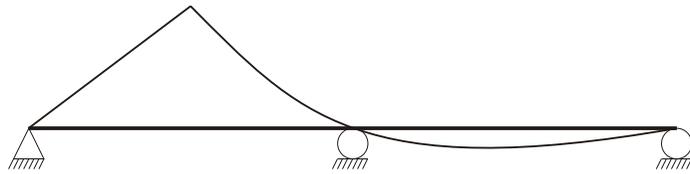
An influence line can be viewed as a curve of influence values plotted at the load points along a traffic Lane. For a given response quantity (force, displacement, or stress) at a given location in the structure, the influence value plotted at a load point is the value of that response quantity due to a unit concentrated downward force acting at that load point. The influence line thus shows the influence upon the given response quantity of a unit force moving along the traffic lane. Figure 86 (page 486) shows some simple examples of influence lines. An influence surface is the extension of this concept into two dimensions across the width of the lane.

Influence lines and surfaces may exhibit discontinuities (jumps) at the location of the response quantity when it is located at a load point on the traffic lane. Discontinuities may also occur where the structure itself is not continuous (e.g., expansion joints).

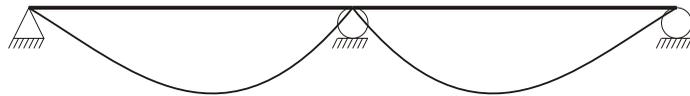
Influence lines and surfaces may be displayed in the graphical user interface for the displacement, force, or stress response of any element in the structure. They are plotted on the Lanes with the influence values plotted in the vertical direction. A positive influence value due to gravity load is plotted upward. Influence values are linearly interpolated between the known values at the load points.



(a) Influence Line for Vertical Shear at Center of a Simple Span



(b) Influence Line for Moment at Center of Left Span of Two Continuous Spans



(c) Influence Line for Moment at Center Support of Two Continuous Spans

Figure 86

Examples of Influence Lines for One-Span and Two-Span Beams

Vehicle Live Loads

Any number of **Vehicle live loads**, or simply **Vehicles**, may be defined to act on the traffic Lanes. You may use standard types of Vehicles known to the program, or design your own using the general Vehicle specification.

Distribution of Loads

Longitudinally, each Vehicle consists of one or more axle loads and/or one or more uniform loads. Axle loads act at a single longitudinal location in the vehicle. Uniform loads may act between pairs of axles, or extend infinitely before the first axle or after the last axle. The width of each axle load and each uniform load is independently specified. These widths may be fixed or equal to the width of the Lane.

For Moving-Load Load Cases using the influence surface, both axle loads and uniform loads are used to maximum effect. For step-by-step analysis, only the axle loads are used.

Axle Loads

Longitudinally, axle loads look like a point load. Transversely, axle loads may be represented as one or more point (wheel) loads or as distributed (knife-edge) loads. Knife-edge loads may be distributed across a fixed width or the full width of the lane. Axle loads may be zero, which can be used to separate uniform loads of different magnitude.

Uniform Loads

Longitudinally, the uniform loads are constant between axles. Leading and trailing loads may be specified that extend to infinity. Transversely, these loads may be distributed uniformly across the width of the lane, over a fixed width, or they may be concentrated at the center line of the lane.

Minimum Edge Distances

Certain AASHTO vehicles require that the wheel loads maintain a specified minimum distance from the edge of the lane. For any vehicle, you may specify a minimum distance for interior edges of lanes, and another distance for exterior edges. By default, these distances are zero. The specified distances apply equally to all axle loads, but do not affect longitudinally-uniform loads. The definition of interior and exterior edges is given in Subtopic “Lanes” (page 482).

Directions of Loading

All vehicle live loads are specified as weight and are assumed, by default, to act downward in the $-Z$ global coordinate direction. However, horizontal loads representing braking, acceleration, and centrifugal forces may also be specified as part of the vehicle definition.

See “Upward and Horizontal Directions” (page 13) in Chapter “Coordinate Systems.”

All vertical and horizontal loads are generated from the same distribution of weights, i.e., the axle and uniform loads as described above. The various directions of load are additive, depending upon scale factors defined in the moving load cases, and are computed simultaneously for each vehicle position.

In the discussion that follows, W is the total weight of an axle load or weight per length of a uniform load at a given longitudinal position within a vehicle. All generated vertical or horizontal loads have the same distribution across the width of the vehicle, except as described for overturning moments below.

Vertical Loads

By default the specified axle and uniform loads act in the gravity direction. You may optionally specify that vertical loads are to be excluded if you want to define a special vehicle that represents horizontal loading only. Accordingly, the vertical force F_V is either

$$F_V = W \text{ or } F_V = 0$$

where W is the axle or uniform weight of the vehicle at a given location.

Note that vertical load is required to generate both super-elevation and centrifugal overturning moments, as discussed below.

Braking and Acceleration Loads

Braking and acceleration loads may be considered by specifying one or two sets of coefficients multiplying the axle and uniform loads. For example, the first set could specify a coefficient of 0.25 on the axle loads and 0.0 on the uniform loads. The second set, if requested, could specify a coefficient of 0.05 on both the axle and the uniform loads. If two sets are requested, the envelope of the two sets will be determined and reported.

The specified values effectively act as coefficients of friction, and the longitudinal force F_B is computed as

$$F_B = bW$$

where b is the specified braking coefficient for axle or uniform loads, and W is the axle or uniform weight of the vehicle at a given location.

Braking and acceleration loads act longitudinally, parallel to the centerline of the Lane, and along the grade of the layout line used to define the Lane. These forces will be assumed to act both forward and backward, so that the maximum and minimum response reported will be equal and opposite to each other. They will be based on either the maximum or minimum response computed, whichever has the larger absolute value.

Optionally, skew braking and acceleration loads can be considered by specifying skew coefficients for axle and uniform loads. The coefficients multiply the longitudinal braking and acceleration forces, and are applied transversely, perpendicular to the Lane centerline. If the Lane is acting on a bridge object, the skew forces will act along the super-elevation slope of the bridge. Otherwise, skew braking forces act in the horizontal plane. Skew forces acting to the left and right, simultaneously with the longitudinal forces, will both be considered and the envelope reported. The skew coefficient is effectively the tangent of the angle between the direction of braking or acceleration and the centerline of the Lane, and the skew braking/acceleration force F_K is computed as

$$F_K = kF_B$$

where k is the skew coefficient.

Centrifugal Loads

Centrifugal loads may be considered by specifying a velocity and a set of coefficients multiplying the axle and uniform loads. The centrifugal force F_C is computed as

$$F_C = c \frac{V_C^2}{g} \frac{W}{R}$$

where c is the specified centrifugal coefficient for axle or uniform loads, V_C is the specified velocity, g is the acceleration due to gravity, W is the axle or uniform weight of the vehicle at a given location, and R is the centrifugal radius specified as part of the Lane definition. When $R = 0$, the radius is taken as infinity (straight) and the force $F_C = 0$.

For the Eurocode and other prescriptive codes, a reference length (typically one meter) may be specified instead of a velocity. In this case the centrifugal force is computed as

$$F_C = cL_C \frac{W}{R}$$

where L_C is the reference length.

Centrifugal forces are applied transversely, perpendicular to the centerline of the Lane, and always act in the horizontal plane regardless of any super-elevation. Note that the force acts to the left (negative) when the radius curves to the right (positive), and vice-versa.

Overtuning Moments

Two types of overturning moments may be considered, centrifugal and super-elevation. Both are proportional to the height of the centroid of vehicle weight above the surface of the bridge deck. Separate heights are specified for axle and uniform loads, but the same heights apply to both types of overturning moment.

For centrifugal loads, the overturning moment (torque) T_C is given by

$$T_C = F_C h$$

where F_C is the centrifugal force as defined above, and h is the height to the centroid of the vehicle weight causing this force. This generates a vertical force couple F_{TC} acting over a moment arm equal to the wheel spacing s , computed as

$$F_{TC} = \frac{T_C}{s} = F_C \frac{h}{s}$$

For vertical loads, the overturning moment T_V is given by

$$T_V = F_V h \sin e$$

where F_V is the vertical force as defined above, h is the height to the centroid of the vehicle weight causing the force, $\tan^{-1} e$ is the angle of super-elevation, and e is the super-elevation defined as a ratio. This generates a vertical force couple F_{TV} acting over a moment arm equal to the horizontal wheel spacing $s \cos e$, computed as

$$F_{TV} = \frac{T_V}{s \cos e} = T_V e \frac{h}{s}$$

The total overturning vertical force couple F_T is given by

$$F_T = F_{TC} - F_{TV} - (F_C - F_V e) \frac{h}{s}$$

Note that the overturning moment due to super-elevation tends to counteract the centrifugal overturning effects on a curve, but may itself cause overturning when acting alone for vehicles that are stationary. Both cases must be considered in an enveloping analysis or design.

For a vehicle load of width w , the spacing s and the application of the vertical force F_T couple depend on the transverse distribution of the load as follows:

- For a single point load, or for a zero-width lane, $s = 0$ and no overturning moments are considered.
- For two point loads, $s = w$, and the overturning load F_T is added to the two vertical loads as an equal-but-opposite pair at the two wheel locations.
- For four equally-spaced point loads, $s = \frac{2}{3} w$, and overturning loads $F_T/2$ are added to the two vertical loads on one side of the vehicle width and subtracted from the loads on the other side.
- For a uniform transverse distribution of load, $s = \frac{1}{2} w$, and a uniform overturning load $F_T/(w/2)$ is added to the vertical load on one half of the vehicle width and subtracted from the other half.
- Note that the width of the vehicle will be reduced, if necessary, to fit within the width of the lane.

Important Note! The overturning load is restricted so as not to cause uplift on the inner or upward side of the vehicle. This restriction is enforced during analysis, and can be stated as

$$F_T \leq \frac{F_V}{2}$$

regardless of the width distribution of loads.

In particular, no overturning moments will be considered for a load case where the vertical load is zero. When defining moving load cases, be sure to combine vertical load with centrifugal load if overturning moments are important. A centrifugal load case by itself will exclude overturning moments, even if later included in a load combination with vertical load cases. See Topic “Moving-Load Load Cases” (page 509) for more information.

Elevation of Horizontal Loads

Horizontal loads generated for influence surfaces or individual vehicles are applied at vertical elevations intended to represent the top surface of the bridge deck. Specifically:

- For spine and solid models of a bridge object, horizontal loads act at the elevation of the top surface, after taking into account the grade and super-elevation of the bridge object.
- For area (shell) models of a bridge object, the loads act at the joints of the area objects representing the top slab. This is typically at the mid-surface of the top slab, and hence the loads act at a half slab-thickness below the top surface. Grade and super-elevation are included in the location of these joints.
- For models not created with the Bridge Modeler, horizontal loads on frame objects act at the elevation of the layout line. Horizontal loads on area objects act at their joints (typically at the mid-surface), and they act at the top-surface joints of solid objects.

For this reason, torsional moments created by transverse loads and vertical moments (M3) caused by longitudinal loads may be slightly different according to the type of model generated to represent the bridge object. The difference is generally small but should be recognized.

Restricting a Vehicle to the Lane Length

When moving a vehicle along the length of the lane, the front of the vehicle starts at one end of the lane, and the vehicle travels forward until the back of the vehicle exits the other end of the lane. This means that all locations of the vehicle are considered, whether fully or partially on the lane

You have the option to specify that a vehicle must remain fully on the lane. This is useful for cranes and similar vehicles that have stops at the end of their rails that prevent them from leaving the lane. This setting only affects influence-surface analysis, not step-by-step analysis where you can explicitly control where the vehicle runs.

Application of Loads to the Influence Surface

The maximum and minimum values of a response quantity are computed using the corresponding influence line or surface. Concentrated loads are multiplied by the influence value at the point of application to obtain the corresponding response;

distributed loads are multiplied by the influence values and integrated over the length and width of application.

By default, each concentrated or distributed load is considered to represent a range of values from zero up to a specified maximum. When computing a response quantity (force or displacement) the maximum value of load is used where it increases the severity of the response, and zero is used where the load would have a relieving effect. Thus the specified load values for a given Vehicle may not always be applied proportionally. This is a conservative approach that accounts for Vehicles that are not fully loaded. Thus the maximum response is always positive (or zero); the minimum response is always negative (or zero).

You may override this conservative behavior as discussed in the next Subtopic, “Option to Allow Reduced Response Severity”.

By way of example, consider the influence line for the moment at the center of the left span shown in Figure 86(b) (page 486). Any axle load or portion of a distributed load that acts on the left span would contribute only to the positive maximum value of the moment response. Loads acting on the right span would not decrease this maximum, but would contribute to the negative minimum value of this moment response.

Option to Allow Reduced Response Severity

You have the option to allow loads to reduce the severity of the response. If you choose this option, all concentrated and uniform loads will be applied at full value on the entire influence surface, regardless of whether or not that load reduces the severity of the response. *This is less conservative than the default method of load application.* The use of this option may be useful for routing special vehicles whose loads are well known. However, for notional loads that represent a distribution or envelope of unknown vehicle loadings, the default method may be more appropriate.

Width Effects

Fixed-width loads will be moved transversely across the width of a Lane for maximum effect if the Lane is wider than the load. If the Lane is narrower than the load, the load will be centered on the Lane and its width reduced to be equal to that of the Lane, keeping the total magnitude of the load unchanged.

The load at each longitudinal location in the vehicle is independently moved across the width of the Lane. This means that the front, back, and middle of the vehicle

may not occupy the same transverse location in the lane when placed for maximum effect.

Length Effects

You may specify that the magnitude of the loading depends on lane length using built-in or user-defined length functions. One function may be used to affect the concentrated (axle) loads, and another function may be used for the distributed loads. These functions act as scale factors on the specified load values.

Concentrated (Axle) Loads

If you specify a length-effect function for the axle loads, all axle loads will be scaled equally by the function, including floating axle loads. Built-in length-effect functions include the AASHTO Standard Impact function and the JTG-D60 Lane load function. You may define your own function.

The intent of this function is to scale the load according to span length. In a given structure, there may not be a constant span length, so the program uses the influence line to determine what span length to use. This may differ for each computed response quantity, and may not always correspond to the obvious span length in the global structure.

For a given response quantity, the maximum point on the influence line is found, and the distance between the zero-crossings on either side of this maximum is taken to be the span length. For the three influence lines of Figure 86 (page 486), this would result in a span length of half the distance between the supports for the shear in (a), and the full distance between the supports for the moments in (b) and (c). For shear near the support, the span length would be essentially the same as the distance between the support.

This approach generally works well for moments and for shear near the supports. A shorter span length is computed for shear near midspan, but here the shear is smaller anyway, so it is not usually of concern.

Distributed Loads

If you specify a length-effect function for the distributed loads, all distributed loads will be scaled equally by the function. Built-in length-effect functions include the AASHTO Standard Impact function and the British HA function. You may define your own function.

The intent of this function is to scale the load according to the loaded length, but not unconservatively. The influence line is used to determine the loaded length for each individual response quantity. Only loaded lengths that increase the severity of the response are considered.

To prevent long lengths of small influence from unconservatively reducing the response, an iterative approach is used where the length considered is progressively increased until the maximum response is computed. Any further increases in length that reduce the response due to decreasing function value are ignored.

On the other hand, the loaded length used for scaling is never reduced below the length of a single span, which is determined separately for each influence line as the distance between zero crossings on either side of the point of maximum influence.

Application of Loads in Multi-Step Analysis

Vehicles can be moved in a multi-step analysis. This can use either Multi-Step Static Load Cases or Time-History Load Cases, the latter of which can be linear or nonlinear. Currently only vertical loading is considered for multi-step analysis.

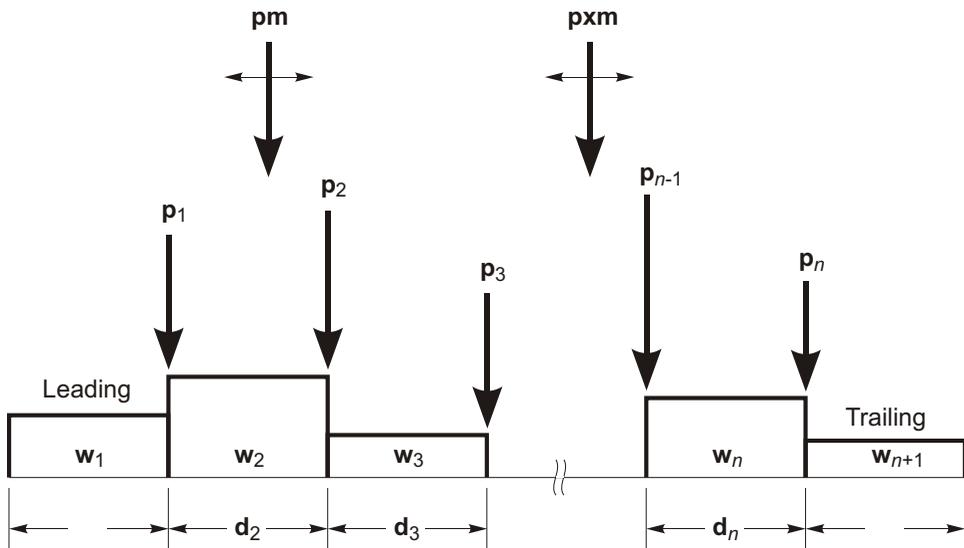
Influence surfaces are not used for this type of analysis. Rather, CSiBridge creates many internal Load Patterns representing different positions of the vehicles along the length of the lanes.

Only axle loads are considered; the uniform loads are not applied. In the case of a variable axle spacing, the minimum distance is used. The transverse distribution of the axle loads is considered. The vehicle is moved longitudinally along the center-line of the lane; it is not moved transversely within the lane. To consider different transverse positions, you can define additional lanes.

The full magnitude of the loads are applied, regardless of whether they increase or decrease the severity of the response. Each step in the analysis corresponds to a specific position of each vehicle acting in its lane. All response at that step is fully correlated.

General Vehicle

The general Vehicle may represent an actual vehicle or a notional vehicle used by a design code. Most trucks and trains can be modeled by the general Vehicle in CSiBridge.



Notes:

- (1) All loads are point loads or uniform line loads acting on the Lane center line
- (2) Any of the point loads or uniform line loads may be zero
- (3) The number of axles, n , may be zero or more
- (4) One of the inter-axle spacings, d_2 through d_n , may vary over a specified range
- (5) The locations of loads pm and pxm are arbitrary

Figure 87
General Vehicle Definition

The general Vehicle consists of n axles with specified distances between them. Concentrated loads may exist at the axles. Uniform loads may exist between pairs of axles, in front of the first axle, and behind the last axle. The distance between any one pair of axles may vary over a specified range; the other distances are fixed. The leading and trailing uniform loads are of infinite extent. Additional “floating” concentrated loads may be specified that are independent of the position of the axles.

By default for influence surface analysis, applied loads never decrease the severity of the computed response, so the effect of a shorter Vehicle is captured by a longer

Vehicle that includes the same loads and spacings as the shorter Vehicle. Only the longer Vehicle need be considered in such cases.

If you choose the option to allow loads to reduce the severity of response, then you must consider the shorter and longer vehicles, if they both apply. This is also true for step-by-step analysis.

Specification

To define a Vehicle, you may specify:

- $n-1$ positive distances, **d**, between the pairs of axles; one inter-axle distance may be specified as a range from **dmin** to **dmax**, where $0 < \mathbf{dmin} \leq \mathbf{dmax}$, and **dmax** = 0 can be used to represent a maximum distance of infinity
- n concentrated loads, **p**, at the axles, including the transverse load distribution for each
- $n+1$ uniform loads, **w**: the leading load, the inter-axle loads, and the trailing load, including the transverse load distribution for each
- Floating axle loads:
 - Load **pm** for superstructure moments, including its transverse distribution. You may specify whether or not to double this load for negative superstructure moments over the supports, as described below
 - Load **pxm** for all response quantities except superstructure moments, including its transverse distribution
- Whether or not this Vehicle is to be used for calculating:
 - “Negative” superstructure moments over the supports
 - Reaction forces at interior supports
 - Response quantities other than the two types above
- Minimum distances between the axle loads and the edges of the lane; by default these distances are zero
- Whether or not the vehicle must remain fully within the length of lane
- Whether or not to automatically scale the magnitude of the axle loads based on the loaded length of the lane according to the Chinese JTG or other codes
- Whether or not to automatically scale the magnitude of the uniform loads based on the loaded length of the lane according to the British or other codes
- Whether or not to consider vertical loading, and whether or not to include the corresponding overturning moments due to super-elevation

- Whether or not to consider braking and acceleration loading, and the scale factors defining these loads
- Whether or not to consider centrifugal loading, the scale factors defining these loads, and whether or not to include the corresponding overturning moments

The number of axles, n , may be zero, in which case only a single uniform load and the floating concentrated loads can be specified.

These parameters are illustrated in Figure 87 (page 496). Some specific examples are given in Topic “Standard Vehicles” (page 500). Additional detail is provided in the following.

Moving the Vehicle

When a Vehicle is applied to a traffic Lane, the axles are moved along the length of the lane to where the maximum and minimum values are produced for every response quantity in every element. Usually this location will be different for each response quantity. For asymmetric (front to back) Vehicles, both directions of travel are considered.

Vehicle Response Components

Certain features of the AASHTO H, HS, and HL vehicular live loads (AASHTO, 2004) apply only to certain types of bridge response, such as negative moment in the superstructure or the reactions at interior supports. CSiBridge uses the concept of **vehicle response components** to identify these response quantities. You select the objects that need special treatment, and assign the appropriate vehicle response components to them.

The different types of available vehicle response components are described in the following subtopics.

Vehicle response components do not apply to SAP2000.

Superstructure (Span) Moment

For AASHTO H and HS “Lane” loads, the floating axle load **pm** is used for calculating the superstructure moment. How this moment is represented depends on the type of model used. For all other types of response, the floating axle load **pxm** is used.

The general procedure is to select the elements representing the superstructure and assign vehicle response components “H and HS Lane Loads – Superstructure Moment” to the desired response quantities, as described next.

For a spine (spline) model where the superstructure is modeled as a line of frame elements, superstructure moment corresponds to frame moment M3 for elements where the local-2 axis is in the vertical plane (the default.) Thus you would select all frame elements representing the superstructure and assign the vehicle response components to M3, indicating to “Use All Values” (i.e., positive and negative.) Load **pm** will be used for computing M3 of these elements.

For a full-shell model of the superstructure, superstructure moment corresponds to longitudinal stresses or membrane forces in the shell elements. Assuming the local-1 axes of the shell elements are oriented along the longitudinal direction of the bridge, you would select all shell elements representing the superstructure and assign the vehicle response components to S11 and/or F11, indicating to “Use All Values” (i.e., positive and negative.) You could also make this same assignment to shell moments M11. Load **pm** will be used for computing any components you have so assigned.

Negative Superstructure (Span) Moment

For AASHTO H and HS “Lane” loads, the floating axle load **pm** is applied in two adjacent spans for calculating the negative superstructure moment over the supports. Similarly, for AASHTO HL loads, a special double-truck vehicle is used for calculating negative superstructure moment over interior supports. Negative moment here means a moment that causes tension in the top of the superstructure, even if the sign of the CSiBridge response is positive due to a particular choice of local axes.

The procedure for different types of structures is very similar to that described above for superstructure moment: select the elements representing the superstructure, but now assign vehicle response components “H, HS and HL Lane Loads – Superstructure Negative Moment over Supports” to the desired response quantities. However, we have to decide how to handle the sign.

There are two general approaches. Let's consider the case of the spine model with frame moment M3 representing superstructure moment:

- (1) You can select the entire superstructure, and assign the vehicle response components to M3, indicating to “Use Negative Values”. Only negative values of M3 will be computed using the double **pm** or double-truck load.

- (2) You can select only that part of the superstructure within a pre-determined negative-moment region, such as between the inflection points under dead load. Assign the vehicle response components to M3, indicating to “Use Negative Values” or “Use All Values.”

The first approach may be slightly more conservative, giving negative moments over a larger region. However, it does not require you to determine a negative-moment region.

The situation with the shell model is more complicated, since negative moments correspond to positive membrane forces and stresses at the top of the superstructure, negative values at the bottom of the superstructure, and changing sign in between. For this reason, approach (2) above may be better: determine a negative-moment region, then assign the vehicle response components to the desired shell stresses, membrane forces, and/or moments, indicating to “Use All Values.” This avoids the problem of sign where it changes through the depth.

Reactions at Interior Supports

For AASHTO HL loads, a special double-truck vehicle is used for calculating the reactions at interior supports. It is up to you to determine what response components you want to be computed for this purpose. Choices could include:

- Vertical upward reactions, or all reactions, for springs and restraints at the base of the columns
- Compressive axial force, or all forces and moments, in the columns
- Compressive axial force, or all forces and moments, in link elements representing bearings
- Bending moments in outriggers at the columns

The procedure is as above for superstructure moment. Select the elements representing the interior supports and assign the vehicle response components “HL – Reactions at Interior Supports” to the desired response quantities. Carefully decide whether you want to use all values, or only negative or positive values. You will have to repeat this process for each type of element that is part of the interior supports: joints, frames, links, shells, and/or solids.

Standard Vehicles

There are many standard vehicles available in CSiBridge to represent vehicular live loads specified in various design codes. More are being added all the time. A few

examples are provided here for illustrative purposes. Only the longitudinal distribution of loading is shown in the figures. Please see the graphical user interface for all available types and further information.

Hn-44 and HSn-44

Vehicles specified with **type** = Hn-44 and **type** = HSn-44 represent the AASHTO standard H and HS Truck Loads, respectively. The **n** in the **type** is an integer scale factor that specifies the nominal weight of the Vehicle in tons. Thus H15-44 is a nominal 15 ton H Truck Load, and HS20-44 is a nominal 20 ton HS Truck Load. These Vehicles are illustrated in Figure 88 (page 503).

The effect of an H Vehicle is included in an HS Vehicle of the same nominal weight. If you are designing for both H and HS Vehicles, only the HS Vehicle is needed.

Hn-44L and HSn-44L

Vehicles specified with **type** = Hn-44L and **type** = HSn-44L represent the AASHTO standard H and HS Lane Loads, respectively. The **n** in the **type** is an integer scale factor that specifies the nominal weight of the Vehicle in tons. Thus H15-44 is a nominal 15 ton H Lane Load, and HS20-44 is a nominal 20 ton HS Lane Load. These Vehicles are illustrated in Figure 88 (page 503). The Hn-44L and HSn-44L Vehicles are identical.

AML

Vehicles specified with **type** = AML represent the AASHTO standard Alternate Military Load. This Vehicle consists of two 24 kip axles spaced 4 ft apart.

HL-93K, HL-93M and HL-93S

Vehicles specified with **type** = HL-93K represent the AASHTO standard HL-93 Load consisting of the code-specified design truck and the design lane load.

Vehicles specified with **type** = HL-93M represent the AASHTO standard HL-93 Load consisting of the code-specified design tandem and the design lane load.

Vehicles specified with **type** = HL-93S represent the AASHTO standard HL-93 Load consisting of two code-specified design trucks and the design lane load, all scaled by 90%. The axle spacing for each truck is fixed at 14 ft. The spacing between the rear axle of the lead truck and the lead axle of the rear truck varies from 50 ft to the length of the Lane. This vehicle is only used for negative superstructure

moment over supports and reactions at interior supports. The response will be zero for all response quantities that do not have the appropriately assigned vehicle response components.

A dynamic load allowance may be specified for each Vehicle using the parameter **im**. This is the additive percentage by which the concentrated truck or tandem axle loads will be increased. The uniform lane load is not affected. Thus if **im** = 33, all concentrated axle loads for the vehicle will be multiplied by the factor 1.33.

These Vehicles are illustrated in Figure 89 (page 504) for **im** = 0.

P5, P7, P9, P11, and P13

Vehicles specified with **type** = P5, **type** = P7, **type** = P9, **type** = P11, and **type** = P13 represent the Caltrans standard Permit Loads. These Vehicles are illustrated in Figure 90 (page 505).

The effect of a shorter Caltrans Permit Load is included in any of the longer Permit Loads. If you are designing for all of these permit loads, only the P13 Vehicle is needed.

Cooper E 80

Vehicles specified with **type** = COOPERE80 represent the AREA standard Cooper E 80 train load. This Vehicle is illustrated in Figure 91 (page 506).

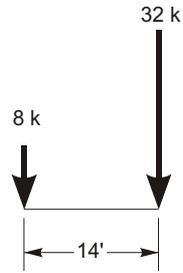
UICn

Vehicles specified with **type** = UICn represent the European UIC (or British RU) train load. The **n** in the **type** is an integer scale factor that specifies magnitude of the uniform load in KN/m. Thus UIC80 is the full UIC load with an 80 kn/m uniform load, and UIC60 is the UIC load with an 60 kn/m uniform load. The concentrated loads are not affected by **n**.

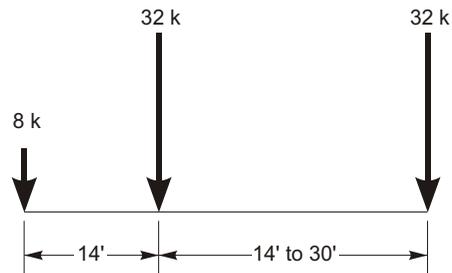
This Vehicle is illustrated in Figure 91 (page 506).

RL

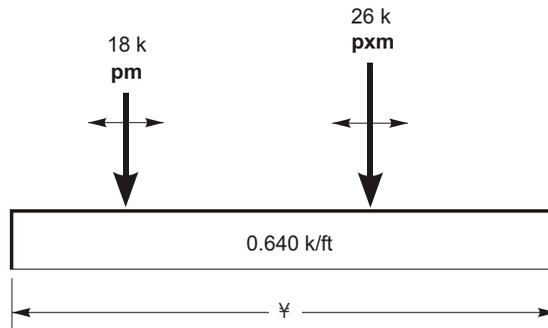
Vehicles specified with **type** = RL represent the British RL train load. This Vehicle is illustrated in Figure 91 (page 506).



H20-44 Truck Load

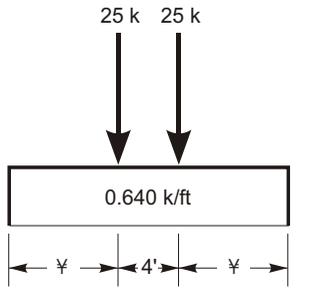


HS20-44 Truck Load



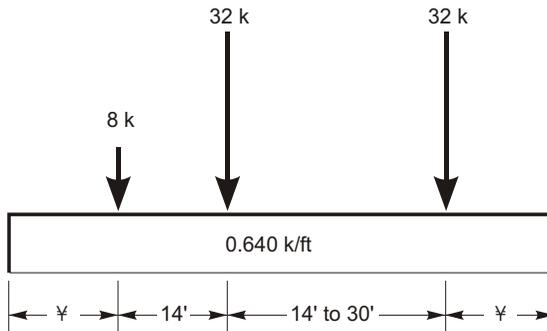
H20-44L and HS20-44L Lane Loads

Figure 88
AASHTO Standard H and HS Vehicles

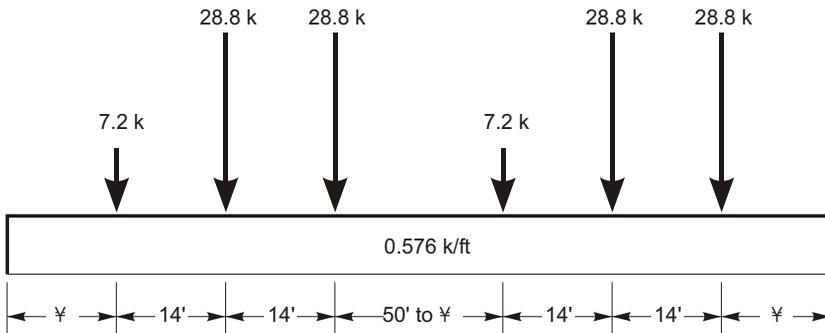


HL-93M Tandem and Lane Load

Note: All point loads will be increased by the dynamic load allowance, im , expressed as a percentage



HL-93K Truck and Lane Load



HL-93S Truck and Lane Load for Negative Moment and Reactions at Interior Piers

Figure 89
AASHTO Standard HL Vehicles

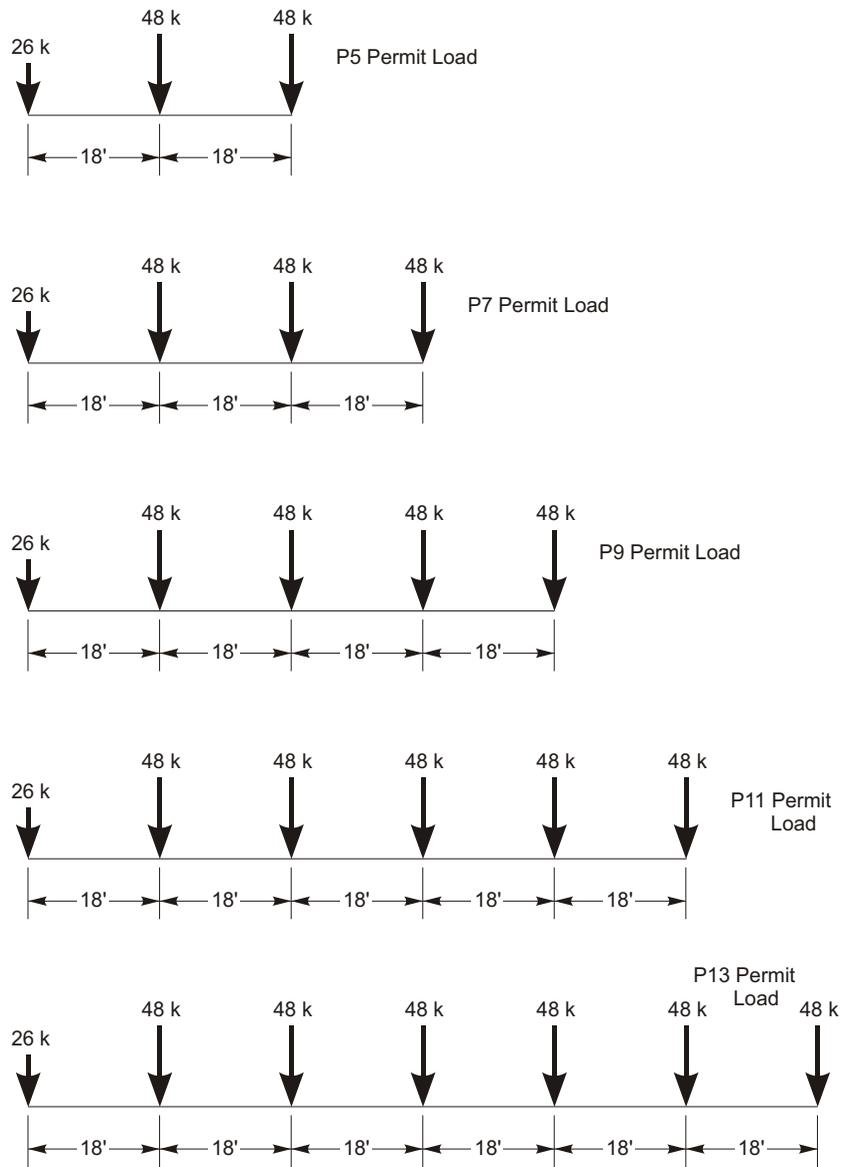
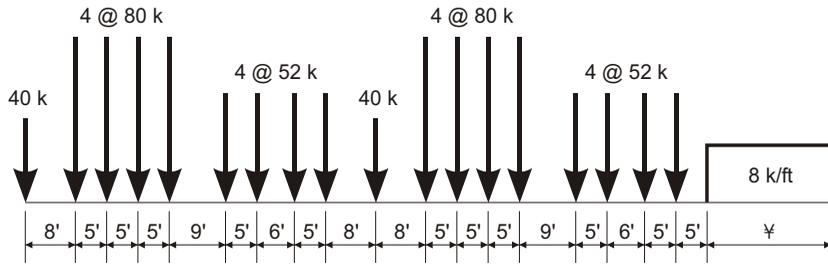
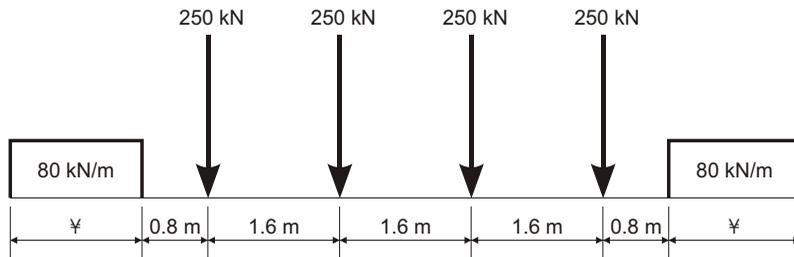


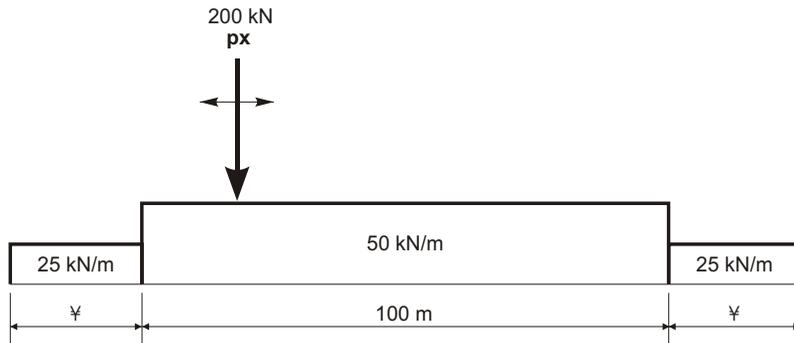
Figure 90
Caltrans Standard Permit Vehicles



Cooper E 80 Train Load



UIC80 Train Load



RL Train Load

Figure 91
Standard Train Vehicles

Eurocode Vehicles

Vehicles are available for the four Load Models of the European Standard EN 1991-2:2003.

There are eight vehicles in total that represent Load Model 1. Each vehicle combines the tandem system TS (concentrated load) with the UDL system (uniform load), but with different scale factors.

For use as a characteristic load, there are separate vehicles for loading the first, second, third, and subsequent lanes. These are given in the table below:

Vehicle	Location	TS, Axle Load Q_{ik} (KN)	UDL Load q_{ik} (KN/m ²)
Euro_LM1_Lane1_Char	Lane Number 1	300	9
Euro_LM1_Lane2_Char	Lane Number 2	200	2.5
Euro_LM1_Lane3_Char	Lane Number 3	100	2.5
Euro_LM1_Other_Char	Other lanes	0	2.5

A moving-load case defined for use as a characteristic load, and having four or more lanes, should apply a single instance each of the first three vehicles to all the lanes and apply multiple instances of the last vehicle to all lanes. The automatic permutations in the moving load case will consider the most severe combination of the vehicles loading the various lanes. If there are fewer than four lanes, the latter vehicles are not needed, respectively.

For use as a frequent load, there are separate vehicles for loading the first, second, third, and subsequent lanes. These are given in the table below:

Vehicle	Location	TS, Axle Load Q_{ik} (KN)	UDL Load q_{ik} (KN/m ²)
Euro_LM1_Lane1_Freq	Lane Number 1	225	3.6
Euro_LM1_Lane2_Freq	Lane Number 2	150	1.0
Euro_LM1_Lane3_Freq	Lane Number 3	75	1.0
Euro_LM1_Other_Freq	Other lanes	0	1.0

Because of the different ratios of TS and UDL load, separate moving-load cases need to be defined for characteristic and frequent loading.

A moving-load case defined for use as a frequent load, and having four or more lanes, should apply a single instance each of the first three vehicles to all the lanes and apply multiple instances of the last vehicle to all lanes. The automatic permutations in the moving load case will consider the most severe combination of the vehicles loading the various lanes. If there are fewer than four lanes, the latter vehicles are not needed, respectively.

A single vehicle is available for Load Model 2, Euro_LM2, that applies a single axle load. This should be applied in a moving-load case that loads all the lanes, but only a single lane at a time.

There are multiple vehicles available to represent the various types defined for Load Model 3. For vehicles with 100 to 200 kn axle-lines, the axle loads are applied as a knife-edge load distributed over a 2.7 m width. For vehicles with 240 kn axle-lines, the axle loads are applied as a knife-edge load distributed over a 4.2 m width. These can be combined with the Load Model 1 vehicles in a moving-load case, loading separate lanes. The 240 kn axle-lines are not automatically applied to two adjacent lanes. Instead, you will need to define a new lane that is double-wide and use that lane in the moving-load case in place of the two adjacent lanes.

A single vehicle is available for Load Model 4, Euro_LM4, that applies a uniform load. This should be applied in a moving-load case that loads all lanes simultaneously. Additional lanes may need to be defined to represent areas that may be loaded by pedestrians but are not otherwise used for traffic.

Vehicle Classes

The designer is often interested in the maximum and minimum response of the bridge to the most extreme of several types of Vehicles rather than the effect of the individual Vehicles. For this purpose, Vehicle Classes are defined that may include any number of individual Vehicles. The maximum and minimum force and displacement response quantities for a Vehicle Class will be the maximum and minimum values obtained for any individual Vehicle in that Class. Only one Vehicle ever acts at a time.

For influence-based analyses, all Vehicle loads are applied to the traffic Lanes through the use of Vehicle Classes. If it is desired to apply an individual Vehicle load, you must use a Vehicle Class that contains only that single Vehicle. Such single-vehicle Classes are automatically created for each Vehicle, and cannot be changed. You need to define your own Vehicle Classes only if you want to envelope multiple Vehicles or change the scale factor on a single Vehicle.

For example, the you may need to consider the most severe of a Truck Load and the corresponding Lane Load, say the HS20-44 and HS20-44L loads. A Vehicle Class can be defined to contain these two Vehicles. Additional Vehicles, such as the Alternate Military Load type AML, could be included in the Class as appropriate. Different members of the Class may cause the most severe response at different locations in the structure.

For HL-93 loading, you would first define three Vehicles, one each of the standard types HL-93K, HL-93M, and HL-93S. You then could define a single Vehicle Class containing all three Vehicles.

For step-by-step analysis, Vehicle loads are applied directly without the use of Classes, since no enveloping is performed.

Moving-Load Load Cases

The final step in the definition of the influence-based vehicle live loading is the application of the Vehicle Classes to the traffic Lanes. This is done by creating independent **Moving-Load Load Cases**.

A Moving-Load Load Case is a type of Load Case. Unlike most other Load Cases, you cannot apply Load Patterns in a Moving-Load Load Case. Instead, each Moving-Load Load Case consists of a set of assignments that specify how the Classes are assigned to the Lanes.

Each **assignment** in a Moving-Load Load Case requires the following data:

- A Vehicle Class, **class**
- A scale factor, **sf**, multiplying the effect of **class** (the default is unity)
- A list, **lanes**, of one or more Lanes in which **class** may act (the default is all Lanes)
- The minimum number, **lmin**, of Lanes **lanes** in which **class** must act (the default is zero)
- The maximum number, **lmax**, of Lanes **lanes** in which **class** may act (the default is all of **lanes**)

The program looks at all of the assignments in a Moving-Load Load Case, and tries every possible permutation of loading the traffic Lanes with Vehicle Classes that is permitted by the assignments. No Lane is ever loaded by more than one Class at a time.

You may specify multiple-lane scale factors, **rf1**, **rf2**, **rf3**, ..., for each Moving-Load Load Case that multiply the effect of each permutation depending upon the number of loaded Lanes. For example, the effect of a permutation that loads two Lanes is multiplied by **rf2**.

The maximum and minimum response quantities for a Moving-Load Load Case will be the maximum and minimum values obtained for any permutation permitted by the assignments. Usually the permutation producing the most severe response will be different for different response quantities.

The concepts of assignment can be clarified with the help of the following examples.

Directions of Loading

For each moving-load load case, three scale factors are specified to multiply and combine the vertical load (**cfv**), braking/acceleration load (**cfb**), and centrifugal load (**cfc**) for all vehicles acting in this load case. By default, **cfv** = 1 and **cfb** = **cfc** = 0. When two or more factors are non-zero, the effects of the different directions of loading for a given vehicle are considered *simultaneously* for each position of the vehicle in the lane. This is more accurate and less over-conservative than considering the three directions of loading separately and then adding them together later in a load combination.

Important Note! The overturning load due to centrifugal load is restricted so as not to cause uplift on the inner or upward side of the vehicle. In particular, no overturning moments will be considered for a load case where the vertical load is zero. For this reason, it is recommended to combine full vertical load with centrifugal load if overturning moments are important. A centrifugal load case by itself will exclude overturning moments, even if later included in a load combination with vertical load cases. See topic “Vehicle Live Loads” (page 487) for more information.

When multiple directions of loading are considered in a moving-load load case, the following permutations are automatically considered during analysis depending upon the type of combination requested by the choice of scale factors:

Combination	Vertical (V)	Braking (B)	Centrif. (C)	Double
V only	cfv	0	0	
B only	0	cfb	0	*
C only	0	0	cfc	

Combination	Vertical (V)	Braking (B)	Centrif. (C)	Double
V + B	cfv	cfb	0	*
	cfv	-cfb	0	*
V + C	cfv	0	cfc	
	cfv	0	0	
B + C	0	cfb	cfc	*
	0	cfb	0	*
	0	-cfb	cfc	*
	0	-cfb	0	*
V + B + C	cfv	cfb	cfc	*
	cfv	cfb	0	*
	cfv	-cfb	cfc	*
	cfv	-cfb	0	*

For vehicles with a non-zero skew factor specified for braking and acceleration loads, the permutations indicated with (*) in the Double column will be doubled to consider both signs of the transverse load due to skew braking/acceleration.

Note that for the “B only” case, the results are automatically enveloped in both directions without any additional computational effort.

Comparing the computational effort to consider all three directions of loading in a single load case versus running them in separate load cases and then adding them together later in load combinations: A single load case requires four permutations, whereas the three individual load cases only require three permutations. However, running all three in a single load case generally provides better results and allows consideration of overturning moments. If the three directions of loading are to be combined with various different sets of scale factors, doing so with load combinations is clearly more efficient than running multiple load cases, although again less accurate.

In general, it is recommended to combine the various directions of loading in the load cases rather than load combinations.

Example I — AASHTO HS Loading

Consider a four-lane bridge designed to carry AASHTO HS20-44 Truck and Lane Loads, and the Alternate Military Load (AASHTO, 1996). Suppose that it is required that the number of Lanes loaded be that which produces the most severe response in every member. Only one of the three Vehicle loads is allowed per lane. Load intensities may be reduced by 10% and 25% when three or four Lanes are loaded, respectively.

Generally, loading all of the Lanes will produce the most severe moments and shears along the span and axial forces in the piers. However, the most severe torsion of the bridge deck and transverse bending of the piers will usually be produced by loading only those Lanes possessing eccentricities of the same sign.

Assume that the bridge structure and traffic Lanes have been defined. Three Vehicles are defined:

- **name** = HSK, **type** = HS20-44
- **name** = HSL, **type** = HS20-44L
- **name** = AML, **type** = AML

where **name** is an arbitrary label assigned to each Vehicle. The three Vehicles are assigned to a single Vehicle Class, with an arbitrary label of **name** = HS, so that the most severe of these three Vehicle loads will be used for every situation.

A single Moving-Load Load Case is then defined that seeks the maximum and minimum responses throughout the structure for the most severe of loading all four Lanes, any three Lanes, any two Lanes or any single Lane. This can be accomplished using a single assignment. The parameters for the assignment are:

- **class** = HS
- **sf** = 1
- **lanes** = 1, 2, 3, 4
- **lmin** = 1
- **lmax** = 4

The scale factors for the loading of multiple Lanes in the set of assignments are **rf1** = 1, **rf2** = 1, **rf3** = 0.9, and **rf4** = 0.75.

There are fifteen possible permutations assigning the single Vehicle Class HS to any one, two, three, or four Lanes. These are presented in the following table:

Permutation	Lane 1	Lane 2	Lane 3	Lane 4	Scale Factor
1	HS				1.00
2		HS			1.00
3			HS		1.00
4				HS	1.00
5	HS	HS			1.00
6		HS	HS		1.00
7			HS	HS	1.00
8	HS			HS	1.00
9	HS		HS		1.00
10		HS		HS	1.00
11	HS	HS	HS		0.90
12	HS	HS		HS	0.90
13	HS		HS	HS	0.90
14		HS	HS	HS	0.90
15	HS	HS	HS	HS	0.75

An “HS” in a Lane column of this table indicates application of Class HS; a blank indicates that the Lane is unloaded. The scale factor for each permutation is determined by the number of Lanes loaded.

Example 2 — AASHTO HL Loading

Consider a four-lane bridge designed to carry AASHTO HL-93 loading (AASHTO, 2004). The approach is the same as used for AASHTO HS loading in the previous example. Only the multiple-lane scale factors and the Vehicles differ.

Three Vehicles are defined:

- **name** = HLK, **type** = HL-93K
- **name** = HLM, **type** = HL-93M
- **name** = HLS, **type** = HL-93S

where **name** is an arbitrary label assigned to each Vehicle.

The three Vehicles are assigned to a single Vehicle Class, with an arbitrary label of **name** = HL, so that the most severe of these three Vehicle loads will be used for every situation. By definition of the standard Vehicle type HL-93S, Vehicle HLS will only be used when computing negative moments over supports or the reaction at interior piers. The other two Vehicles will be considered for all response quantities.

A single Moving-Load Load Case is then defined that is identical to that of the previous example, except that **class** = HL, and the scale factors for multiple Lanes are **rf1** = 1.2, **rf2** = 1, **rf3** = 0.85, and **rf4** = 0.65.

There are again fifteen possible permutations assigning the single Vehicle Class HL to any one, two, three, or four Lanes. These are similar to the permutations of the previous example, with the scale factors changed as appropriate.

Example 3 — Caltrans Permit Loading

Consider the four-lane bridge of the previous examples now subject to Caltrans Combination Group I_{PW} (Caltrans, 1995). Here the permit load(s) are to be used alone in a single traffic Lane, or in combination with one HS or Alternate Military Load in a separate traffic lane, depending upon which is more severe.

Four Vehicles are defined:

- **name** = HSK, **type** = HS20-44
- **name** = HSL, **type** = HS20-44L
- **name** = AML, **type** = AML
- **name** = P13, **type** = P13

where **name** is an arbitrary label assigned to each Vehicle.

The first three Vehicles are assigned to a Vehicle Class that is given the label **name** = HS, as in Example 1. The last Vehicle is assigned as the only member of a Vehicle Class that is given the label **name** = P13. Note that the effects of CSiBridge Vehicle types P5, P7, P9, and P11 are captured by Vehicle type P13.

Combination Group I_{PW} is then represented as a single Moving-Load Load Case consisting of the assignment of Class P13 to any single Lane with or without Class HS being assigned to any other single Lane. This can be accomplished using two assignments. A scale factor of unity is used regardless of the number of loaded Lanes.

The first assignment assigns Class P13 to any single Lane:

- **class** = P13
- **sf** = 1
- **lanes** = 1, 2, 3, 4
- **lmin** = 1
- **lmax** = 1

The second assignment assigns Class HS to any single Lane, or to no Lane at all:

- **class** = HS
- **sf** = 1
- **lanes** = 1, 2, 3, 4
- **lmin** = 0
- **lmax** = 1

There are sixteen possible permutations for these two assignments such that no Lane is loaded by more than one Class at a time. These are presented in the following table:

Permutation	Lane 1	Lane 2	Lane 3	Lane 4	Scale Factor
1	P				1.00
2	P	HS			1.00
3	P		HS		1.00
4	P			HS	1.00
5	HS	P			1.00
6		P			1.00
7		P	HS		1.00
8		P		HS	1.00
9	HS		P		1.00
10		HS	P		1.00
11			P		1.00

Permutation	Lane 1	Lane 2	Lane 3	Lane 4	Scale Factor
12			P	HS	1.00
13	HS			P	1.00
14		HS		P	1.00
15			HS	P	1.00
16				P	1.00

Example 4 — Restricted Caltrans Permit Loading

Consider the four-Lane bridge and the Caltrans permit loading of Example 3, but subject to the following restrictions:

- The permit Vehicle is only allowed in Lane 1 or Lane 4
- The Lane adjacent to the Lane occupied by the permit Vehicle must be empty

Two Moving-Load Load Cases are required, each containing two assignments. A scale factor of unity is used regardless of the number of loaded Lanes.

The first Moving-Load Load Case considers the case where the permit Vehicle occupies Lane 1. The first assignment assigns Class P13 to Lane 1

- **class** = P13
- **sf** = 1
- **lanes** = 1
- **lmin** = 1
- **lmax** = 1

The second assignment assigns Class HS to either Lane 3 or 4, or to no Lane at all:

- **class** = HS
- **sf** = 1
- **lanes** = 3, 4
- **lmin** = 0
- **lmax** = 1

These assignments permits the following three permutations:

Permutation	Lane 1	Lane 2	Lane 3	Lane 4	Scale Factor
1	P				1.00
2	P		HS		1.00
3	P			HS	1.00

Similarly, the second Moving-Load Load Case considers the case where the permit Vehicle occupies Lane 4. The first assignment assigns Class P13 to Lane 4

- **class** = P13
- **sf** = 1
- **lanes** = 4
- **lmin** = 1
- **lmax** = 1

The second assignment assigns Class HS to either Lane 1 or 2, or to no Lane at all:

- **class** = HS
- **sf** = 1
- **lanes** = 1, 2
- **lmin** = 0
- **lmax** = 1

These assignments permits the following three permutations:

Permutation	Lane 1	Lane 2	Lane 3	Lane 4	Scale Factor
1				P	1.00
2	HS			P	1.00
3		HS		P	1.00

An envelope-type Combo that includes only these two Moving-Load Load Cases would produce the most severe response for the six permutations above.

See Topic “Load Combinations (Combos)” (page 351) in Chapter “Load Cases” for more information.

Example 5 — Eurocode Characteristic Load Model I

Consider a five-lane bridge Eurocode Load Model 1 for use as a characteristic load. There are separate vehicles for each of the first three lanes, and a fourth vehicle for the remaining lanes. All permutations of vehicles loading the lanes need to be considered.

Four Vehicles are defined:

- **name** = Load1, **type** = Euro_LM1_Lane1_Char
- **name** = Load2, **type** = Euro_LM1_Lane2_Char
- **name** = Load3, **type** = Euro_LM1_Lane3_Char
- **name** = LoadX, **type** = Euro_LM1_Other_Char

where **name** is an arbitrary label assigned to each Vehicle.

Each vehicle is assigned automatically to a Vehicle Class of the same name.

The characteristic load is then represented as a single Moving-Load Load Case consisting of the assignment of each of the first three vehicles to a single lane, and the fourth vehicle to the remaining lanes. This can be accomplished using four assignments. A scale factor of unity is used regardless of the number of loaded Lanes.

The first assignment assigns Class Load 1 to any single Lane:

- **class** = Load1
- **sf** = 1
- **lanes** = 1, 2, 3, 4, 5
- **lmin** = 1
- **lmax** = 1

A similar assignment is made for the second and third vehicles.

The fourth assignment assigns Class LoadX to the remaining lanes:

- **class** = LoadX
- **sf** = 1
- **lanes** = 1, 2, 3, 4, 5
- **lmin** = 2
- **lmax** = 2

There are 60 possible permutations for these four assignments such that every Lane is loaded by a single Class at a time. These permutations will be automatically created. Some of these are presented in the following table:

Permutation	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5
1	Load1	Load2	Load3	LoadX	LoadX
2	Load1	Load2	LoadX	Load3	LoadX
3	Load1	Load2	LoadX	LoadX	Load3
4	Load1	Load3	Load2	LoadX	LoadX
5	Load1	Load3	LoadX	Load2	LoadX
...
60	LoadX	LoadX	Load3	Load2	Load1

Considering the possibility that the lanes loaded by LoadX are not loaded at all increases the number of permutations from 60 to 240. However, this is not necessary, since the loading will not be applied in the case that it decreases the severity of the response.

A similar moving-load case can be defined for applying the frequent Load Model 1 vehicles to the Lanes.

Moving Load Response Control

Several parameters are available for controlling influence-based Moving-Load Load Cases. These have no effect on step-by-step analysis.

Bridge Response Groups

By default, moving-load response is calculated for all joints and elements. Because this calculation is computationally intensive, you have the option to restrict this calculation to a portion of the structure.

For each of the following types of response, you may request a Group of elements for which the response should be calculated:

- Joint displacements
- Joint reactions

- Frame forces and moments
- Shell stresses
- Shell resultant forces and moments
- Plane stresses
- Solid stresses
- Link/support forces and deformations

If the displacements, reactions, spring forces, or internal forces are not calculated for a given joint or Frame element, no Moving Load response can be printed or plotted for that joint or element. Likewise, no response can be printed or plotted for any Combo that contains a Moving-Load Load Case.

Additional control is available as described in the following subtopics.

Correspondence

For each maximum or minimum Frame-element response quantity computed, the corresponding values for the other five internal force and moment components may be determined. For example, the shear, moment, and torque that occur at the same time as the maximum axial force in a Frame element may be computed.

Similarly, corresponding displacements, stresses, forces and moments can be computed for any response quantity of any element type. Only the corresponding values for each joint or element are computed at each output location. If you want to see the full corresponding state of the structure, you must use step-by-step analysis.

By default, no corresponding quantities are computed since this significantly increases the computation time for moving-load response.

Influence Line Tolerance

CSiBridge simplifies the influence lines used for response calculation in order to increase efficiency. A relative tolerance is used to reduce the number of load points by removing those that are approximately duplicated or that can be approximately linearly-interpolated. The default value of this tolerance permits response errors on the order of 0.01%. Setting the tolerance to zero will provide exact results to within the resolution of the analysis.

Exact and Quick Response Calculation

For the purpose of moving a Vehicle along a lane, each axle is placed on every load point in turn. When another axle falls between two load points, the effect of that axle is determined by linear interpolation of the influence values. The effect of uniform loads is computed by integrating the linearly-interpolated segments of the influence line. This method is exact to within the resolution of the analysis, but is computationally intensive if there are many load points.

A “Quick” method is available which may be much faster than the usual “Exact” method, but it may also be less accurate. The Quick method approximates the influence line by using a limited number of load points in each “span.” For purposes of this discussion, a span is considered to be a region where the influence line is all positive or all negative.

The degree of approximation to be used is specified by the parameter **quick**, which may be any non-negative integer. The default value is **quick** = 0, which indicates to use the full influence line, i.e., the Exact method.

Positive values indicate increasing degrees of refinement for the Quick method. For **quick** = 1, the influence line is simplified by using only the maximum or minimum value in each span, plus the zero points at each end of the span. For **quick** = 2, an additional load point is used on either side of the maximum/minimum. Higher degrees of refinement use additional load points. The number of points used in a span can be as many as $2^{\text{quick}}+1$, but not more than the number of load points available in the span for the Exact method.

It is strongly recommended that **quick** = 0 be used for all final analyses. For preliminary analyses, **quick** = 1, 2, or 3 is usually adequate, with **quick** = 2 often providing a good balance between speed and accuracy. The effect of parameter **quick** upon speed and accuracy is problem-dependent, and you should experiment to determine the best value to use for each different model.

Step-By-Step Analysis

Step-by-step analysis can consider any combination of Vehicles operating on the Lanes. Multiple Vehicles can operate simultaneously, even in the same Lane if desired. You define a Load Pattern of type “Vehicle Live,” in which you specify one or more sets of the following:

- Vehicle type
- Lane in which it is traveling

- Starting position in the Lane
- Starting time
- Vehicle speed
- Direction (forward or backward, relative to the Lane direction)

You also specify a time-step size and the total number of time steps to be considered. The total duration of loading is the product of these two. To get a finer spatial discretization of loading, use smaller time steps, or reduce the speed of the vehicles.

Loading

This type of Load Pattern is multi-stepped. It automatically creates a different pattern of loading for each time step. At each step, the load applied to the structure is determined as follows:

- The longitudinal position of each Vehicle in its Lane at the current time is determined from its starting position, speed and direction.
- The Vehicle is centered transversely in the Lane.
- Axle loads are applied to the bridge deck. Concentrated axle loads are applied as specified. Distributed axle loads are converted to four equivalent concentrated loads.
- For each individual concentrated load, consistent joint loads are calculated at the corners of any loaded shell or solid element on the deck. In a spine model, a concentrated force and eccentric moment is applied to the closest frame element representing the superstructure.
- Variable axle spacing, if present, is fixed at the minimum distance.
- Longitudinally-uniform loads are not considered.
- Floating axle loads are not considered.

If you wish to consider different axle spacing, define additional Vehicles. If you wish to consider different transverse placement of the Vehicles, define additional Lanes.

Static Analysis

When a Load Pattern of type “Vehicle Live” is applied in a Multi-Step Static Load Case, it results in a separate linear static solution step for each time step, starting at time zero. Each solution is independent, representing the displacement and stress state in the structure for the current position of the vehicles. You can plot these re-

sults in sequence, create a video showing the movement of the vehicles across the structure along with the resulting displacements and/or stresses, or envelope the results for the Load Case.

Since the analysis is static, the speed of the Vehicles has no effect on the results, other than determining the change in position from one load step to the next.

If any additional loads are assigned to structural objects in Load Pattern of type “Vehicle Live”, these loads are applied equally to the structure in every load step of a multi-step static load case, and act in addition to the vehicle loads.

Time-History Analysis

When a Load Pattern of type “Vehicle Live” is applied in a Time-History Load Case, a separate time function is automatically created for each load pattern that was generated for the vehicle. Each time function ramps the load up from zero to one over one time step, and back down to zero in the succeeding time step. This is done regardless of what time function you may specify with the load pattern in the load case. Thus at any given time within a time step, the applied load due to a Vehicle is a linear interpolation of the load pattern for the vehicle at the beginning and the end of the time step.

If any additional loads are assigned to structural objects in a Load Pattern of type “Vehicle Live”, these loads are applied using the time-function specified with the load pattern, and act in addition to the vehicle loads.

Direct integration is recommended. Modal superposition would require a very large number of modes since the spatial distribution of the load is constantly changing.

Dynamical effects are important in a time-history analysis, and different results may be expected depending upon the speed of the vehicle.

The Time-History Load Case may be linear or nonlinear. If you wish to consider static nonlinearity, you should perform a quasi-static nonlinear time-history analysis, i.e., at very slow speed with long time steps. The speed should be slow enough so that the time it takes to cross a span is significantly longer than the first period of the structure.

Enveloping and Load Combinations

Results for each step-by-step Load Case may be displayed or printed for individual steps, or as an envelope giving the maximum and minimum response. When included in Load Combinations, envelope results will be used.

You can approximate an influence-based analysis by the following technique:

- Define one or more Load Patterns of type Vehicle Live, each of which moves a single Vehicle in a single Lane in a single direction
- For each Load Pattern, create a corresponding Multi-Step Static Load Case that applies only that Load Pattern
- For each Lane, define an envelope-type Load Combination of all Load Cases defined for that Lane
- Define a single range-type Combo that includes all of the Lane envelope-type Load Combinations

You can modify this procedure as needed for your particular application. The important thing is to be sure that in the final Load Combination, no Lane is ever loaded by more than one Vehicle at a time, unless that is your intention.

Influence-based analysis is still more comprehensive, since it includes distributed loads, transverse placement of the Vehicles in the Lanes, variable axle spacing, and more accurate placement of the Vehicles for maximum effect.

See Topic “Load Combinations (Combos)” (page 351) in Chapter “Load Cases” for more information.

Computational Considerations

The computation of influence lines requires a moderate amount of computer time and a large amount of disk storage compared with the execution of other typical CSiBridge analyses. The amount of computer time is approximately proportional to N^2L , where N is the number of structure degrees-of-freedom, and L is the number of load points. The amount of disk storage required is approximately proportional to NL .

The computation of Moving Load response may require a large amount of computer time compared with the execution of other typical CSiBridge analyses. The amount of disk storage needed (beyond the influence lines) is small.

The computation time for Moving Load response is proportional to the number of response quantities requested. The computation time for Moving Load response is also directly proportional to the number of Lanes.

For each Vehicle load, the computation time is approximately proportional to the square of the number of axles. It is also proportional to L , the effective number of load points. Larger values of the truck influence tolerance tend to produce smaller values of L compared to L . The value of L will be different for each response quantity; it tends to be smaller for structures with simple spans than with continuous spans.

For step-by-step analysis, computational time is primarily affected by the number of time steps used. Discretization of the Lanes, and the number and type of Vehicles used has a secondary effect.

Chapter XXVII

References

AASHTO, 2002

Standard Specifications for Highways Bridges, 17th Edition, The American Association of State Highway and Transportation Officials, Inc., Washington, D.C.

AASHTO, 2008

AASHTO LRFD Bridge Design Specifications, 4th Edition 2007, with 2008 Interim Revisions, The American Association of State Highway and Transportation Officials, Inc., Washington, D.C.

ACI, 2005

Building Code Requirements for Structural Concrete (ACI 318-05) and Commentary (ACI 318R-05), American Concrete Institute, Farmington Hills, Mich.

AISC, 2005

ANSI/AISC 360-05: An American National Standard Specification for Structural Steel Buildings, American Institute of Steel Construction, Chicago, Ill.

K. J. Bathe, 1982

Finite Element Procedures in Engineering Analysis, Prentice-Hall, Englewood Cliffs, N.J.

K. J. Bathe and E. L. Wilson, 1976

Numerical Methods in Finite Element Analysis, Prentice-Hall, Englewood Cliffs, N.J.

K. J. Bathe, E. L. Wilson, and F. E. Peterson, 1974

SAP IV — A Structural Analysis Program for Static and Dynamic Response of Linear Systems, Report No. EERC 73-11, Earthquake Engineering Research Center, University of California, Berkeley.

J. L. Batoz and M. B. Tahar, 1982

“Evaluation of a New Quadrilateral Thin Plate Bending Element,” *International Journal for Numerical Methods in Engineering*, Vol. 18, pp. 1655–1677.

Caltrans, 1995

Bridge Design Specifications Manual, as amended to December 31, 1995, State of California, Department of Transportation, Sacramento, Calif.

Comite Euro-International Du Beton, 1993

CEB-FIP Modal Code, Thomas Telford, London

P. C. Roussis and M. C. Constantinou, 2005

Experimental and Analytical Studies of Structures Seismically Isolated with and Uplift-Restraint Isolation System, Report No. MCEER-05-0001, MCEER, State University of New York, Buffalo

R. D. Cook, D. S. Malkus, and M. E. Plesha, 1989

Concepts and Applications of Finite Element Analysis, 3rd Edition, John Wiley & Sons, New York, N.Y.

R. D. Cook and W. C. Young, 1985

Advanced Mechanics of Materials, Macmillan, New York, N.Y.

R. K. Dowell, F. S. Seible, and E. L. Wilson, 1998

“Pivot Hysteretic Model for Reinforced Concrete Members,” *ACI Structural Journal*, Vol. 95, pp. 607–617.

FEMA, 2000

Prestandard and Commentary for Seismic Rehabilitation of Buildings, Prepared by the American Society of Civil Engineers for the Federal Emergency Management Agency (Report No. FEMA-356), Washington, D.C.

A. K. Gupta, 1990

Response Spectrum Method in Seismic Analysis and Design of Structures, Blackwell Scientific Publications, Cambridge, Mass.

J. P. Hollings and E. L. Wilson, 1977

3–9 Node Isoparametric Planar or Axisymmetric Finite Element, Report No. UC SESM 78-3, Division of Structural Engineering and Structural Mechanics, University of California, Berkeley.

T. J. R. Hughes, 2000

The Finite Element Method: Linear Static and Dynamic Finite Element Analysis, Dover.

A. Ibrahimbegovic and E. L. Wilson, 1989

“Simple Numerical Algorithms for the Mode Superposition Analysis of Linear Structural Systems with Non-proportional Damping,” *Computers and Structures*, Vol. 33, No. 2, pp. 523–531.

A. Ibrahimbegovic and E. L. Wilson, 1991

“A Unified Formulation for Triangular and Quadrilateral Flat Shell Finite Elements with Six Nodal Degrees of Freedom,” *Communications in Applied Numerical Methods*, Vol. 7, pp. 1–9.

K. Kasai, O. Takahashi, and Y. Sekiguchi, 2004

“JSSI Manual for Building Passive Control Technology, Part-10 Time-History Analysis Model for Nonlinear Oil Dampers,” Paper 1283, *Proceedings of The 13th World Conference on Earthquake Engineering*, Vancouver, Canada.

M. A. Ketchum, 1986

Redistribution of Stresses in Segmentally Erected Prestressed Concrete Bridges, Report No. UCB/SESM-86/07, Department of Civil Engineering, University of California, Berkeley.

N. Makris and J. Zhang, 2000

“Time-domain Viscoelastic Analysis of Earth Structures,” *Earthquake Engineering and Structural Dynamics*, Vol. 29, pp. 745–768.

L. E. Malvern, 1969

Introduction to the Mechanics of a Continuous Medium, Prentice-Hall, Englewood Cliffs, N.J.

C. Menun and A. Der Kiureghian, 1998

“A Replacement for the 30%, 40%, and SRSS Rules for Multicomponent Seismic Analysis,” *Earthquake Spectra*, Vol. 14, No. 1, pp. 153-163.

T. A. Morgan and S. A. Mahin, 2008

“The Optimization of Multi-Stage Friction Pendulum Isolators for Loss Mitigation Considering a Range of Seismic Hazard,” *Proceedings of The 14th World Conference on Earthquake Engineering*, Beijing, China

S. Nagarajaiah, A. M. Reinhorn, and M. C. Constantinou, 1991

3D-Basis: Nonlinear Dynamic Analysis of Three-Dimensional Base Isolated Structures: Part II, Technical Report NCEER-91-0005, National Center for Earthquake Engineering Research, State University of New York at Buffalo, Buffalo, N. Y.

NRC, 2006

“Combining Modal Responses and Spatial Components in Seismic Response Analysis,” Regulatory Guide 1.92, Revision 2, U.S. Nuclear Regulatory Commission.

Y. J. Park, Y. K. Wen, and A. H-S. Ang, 1986

“Random Vibration of Hysteretic Systems under Bi-Directional Ground Motions,” *Earthquake Engineering and Structural Dynamics*, Vol. 14.

R. J. Roark and W. C. Young, 1975

- Formulas for Stress and Strain*. 5th Edition, McGraw-Hill, New York, N.Y.
- T. Takeda, M. A. Sozen, and N. N. Nielsen, 1970
“Reinforced Concrete Response to Simulated Earthquakes,” *J. Struct. Engrg. Div.*, ASCE, Vol. 96, No. 12, pp. 2257–2273.
- R. L. Taylor and J. C. Simo, 1985
“Bending and Membrane Elements for Analysis of Thick and Thin Shells,” *Proceedings of the NUMETA 1985 Conference*, Swansea, Wales.
- K. Terzaghi and R. B. Peck, 1967
Soil Mechanics in Engineering Practice, 2nd Edition, John Wiley & Sons, New York, N.Y.
- S. Timoshenko and S. Woinowsky-Krieger, 1959
Theory of Plates and Shells, 2nd Edition, McGraw-Hill, New York, N.Y.
- Y. K. Wen, 1976
“Method for Random Vibration of Hysteretic Systems,” *Journal of the Engineering Mechanics Division*, ASCE, Vol. 102, No. EM2.
- D. W. White and J. F. Hajjar, 1991
“Application of Second-Order Elastic Analysis in LRFD: Research to Practice,” *Engineering Journal*, AISC, Vol. 28, No. 4, pp. 133–148.
- E. L. Wilson, 1970
SAP — A General Structural Analysis Program, Report No. UC SESM 70-20, Structural Engineering Laboratory, University of California, Berkeley.
- E. L. Wilson, 1972
SOLID SAP — A Static Analysis Program for Three Dimensional Solid Structures, Report No. UC SESM 71-19, Structural Engineering Laboratory, University of California, Berkeley.
- E. L. Wilson, 1985
“A New Method of Dynamic Analysis for Linear and Non-Linear Systems,” *Finite Elements in Analysis and Design*, Vol. 1, pp. 21–23.
- E. L. Wilson, 1993

“An Efficient Computational Method for the Base Isolation and Energy Dissipation Analysis of Structural Systems,” ATC17-1, *Proceedings of the Seminar on Seismic Isolation, Passive Energy Dissipation, and Active Control*, Applied Technology Council, Redwood City, Calif.

E. L. Wilson, 1997

Three Dimensional Dynamic Analysis of Structures with Emphasis on Earthquake Engineering, Computers and Structures, Inc., Berkeley, Calif.

E. L. Wilson and M. R. Button, 1982

“Three Dimensional Dynamic Analysis for Multicomponent Earthquake Spectra,” *Earthquake Engineering and Structural Dynamics*, Vol. 10.

E. L. Wilson, A. Der Kiureghian, and E. P. Bayo, 1981

“A Replacement for the SRSS Method in Seismic Analysis,” *Earthquake Engineering and Structural Dynamics*, Vol. 9.

E. L. Wilson and I. J. Tetsuji, 1983

“An Eigensolution Strategy for Large Systems,” *Computers and Structures*, Vol. 16.

E. L. Wilson, M. W. Yuan, and J. M. Dickens, 1982

“Dynamic Analysis by Direct Superposition of Ritz Vectors,” *Earthquake Engineering and Structural Dynamics*, Vol. 10, pp. 813–823.

V. Zayas and S. Low, 1990

“A Simple Pendulum Technique for Achieving Seismic Isolation,” *Earthquake Spectra*, Vol. 6, No. 2.

O. C. Zienkiewicz and R. L. Taylor, 1989

The Finite Element Method, 4th Edition, Vol. 1, McGraw-Hill, London.

O. C. Zienkiewicz and R. L. Taylor, 1991

The Finite Element Method, 4th Edition, Vol. 2, McGraw-Hill, London.

Copyright Notice for TAUCS:

TAUCS Version 2.0, November 29, 2001. Copyright ©) 2001, 2002, 2003 by Sivan Toledo, Tel-Aviv University, stoledo@tau.ac.il. All Rights Reserved.

TAUCS License:

Your use or distribution of TAUCS or any derivative code implies that you agree to this License.

THIS MATERIAL IS PROVIDED AS IS, WITH ABSOLUTELY NO WARRANTY EXPRESSED OR IMPLIED. ANY USE IS AT YOUR OWN RISK.

Permission is hereby granted to use or copy this program, provided that the Copyright, this License, and the Availability of the original version is retained on all copies. User documentation of any code that uses this code or any derivative code must cite the Copyright, this License, the Availability note, and “Used by permission.” If this code or any derivative code is accessible from within MATLAB, then typing “help taucs” must cite the Copyright, and “type taucs” must also cite this License and the Availability note. Permission to modify the code and to distribute modified code is granted, provided the Copyright, this License, and the Availability note are retained, and a notice that the code was modified is included. This software is provided to you free of charge.

Availability:

As of version 2.1, we distribute the code in 4 formats: zip and tarred-gzipped (tgz), with or without binaries for external libraries. The bundled external libraries should allow you to build the test programs on Linux, Windows, and MacOS X without installing additional software. We recommend that you download the full distributions, and then perhaps replace the bundled libraries by higher performance ones (e.g., with a BLAS library that is specifically optimized for your machine). If you want to conserve bandwidth and you want to install the required libraries yourself, download the lean distributions. The zip and tgz files are identical, except that on Linux, Unix, and MacOS, unpacking the tgz file ensures that the configure script is marked as executable (unpack with `tar xzvpf`), otherwise you will have to change its permissions manually.

Used by permission.

