

MBDyn Input File Format

Version 1.4.2

Pierangelo Masarati

DIPARTIMENTO DI INGEGNERIA AEROSPAZIALE
POLITECNICO DI MILANO

Automatically generated January 8, 2012

Contents

1	Overview	13
1.1	Execution	13
1.1.1	Passing the Input	13
1.1.2	Passing the Output File Name	13
1.2	Input File Structure	14
1.3	Output	15
2	General	17
2.1	Types	17
2.1.1	Keywords	17
2.1.2	Strings	17
2.1.3	Numeric Values	18
2.1.4	Namespaces	18
2.1.5	Plugin Variables	20
2.2	Higher-Order Math Structures	25
2.2.1	3×1 vectors	25
2.2.2	6×1 vectors	26
2.2.3	3×3 matrices	26
2.2.4	3×3 orientation matrices	27
2.2.5	6×6 matrices	29
2.2.6	$6 \times N$ matrices	29
2.3	Input Related Cards	29
2.3.1	Constitutive Law	30
2.3.2	C81 Data	30
2.3.3	Drive Caller	32
2.3.4	Hydraulic fluid	32
2.3.5	Include	33
2.3.6	Module Load	33
2.3.7	Print symbol table	34
2.3.8	Reference	34
2.3.9	Remark	38
2.3.10	Set	38
2.3.11	Setenv	39
2.3.12	Template Drive Caller	39
2.4	Node Degrees of Freedom	40
2.5	Drives and Drive Callers	40
2.5.1	Array drive	41
2.5.2	Closest next drive	41

2.5.3	Const(ant) drive	41
2.5.4	Cosine drive	41
2.5.5	Cubic drive	43
2.5.6	Direct	43
2.5.7	Dof drive	44
2.5.8	Double ramp drive	44
2.5.9	Double step drive	45
2.5.10	Drive drive	45
2.5.11	Element drive	46
2.5.12	Exponential drive	46
2.5.13	File drive	47
2.5.14	Fourier series drive	47
2.5.15	Frequency sweep drive	48
2.5.16	GiNaC	48
2.5.17	Linear drive	49
2.5.18	Meter drive	49
2.5.19	Mult drive	49
2.5.20	Node drive	49
2.5.21	Null drive	50
2.5.22	Parabolic drive	50
2.5.23	Periodic drive	50
2.5.24	Piecewise linear drive	50
2.5.25	Ramp drive	51
2.5.26	Random drive	51
2.5.27	Sample and hold drive	51
2.5.28	Scalar function drive	51
2.5.29	Sine drive	52
2.5.30	Step drive	53
2.5.31	String drive	53
2.5.32	Tanh drive	53
2.5.33	Time drive	53
2.5.34	Timestep drive	54
2.5.35	Unit drive	54
2.5.36	Deprecated drive callers	54
2.5.37	Hints	54
2.5.38	Template Drive	55
2.6	Scalar functions	56
2.7	Friction	61
2.7.1	Friction models	61
2.7.2	Shape functions	62
2.8	Shapes	63
2.9	Constitutive Laws	64
2.9.1	Linear elastic, linear elastic isotropic	66
2.9.2	Linear elastic generic	66
2.9.3	Linear elastic generic axial torsion coupling	67
2.9.4	Cubic elastic generic	67
2.9.5	Log elastic	67
2.9.6	Linear elastic bistop	67
2.9.7	Double linear elastic	68

2.9.8	Isotropic hardening elastic	68
2.9.9	Scalar function elastic, scalar function elastic isotropic	68
2.9.10	Scalar function elastic orthotropic	68
2.9.11	Linear viscous, linear viscous isotropic	69
2.9.12	Linear viscous generic	69
2.9.13	Linear viscoelastic, linear viscoelastic isotropic	69
2.9.14	Linear viscoelastic generic	70
2.9.15	Linear time variant viscoelastic generic	70
2.9.16	Linear viscoelastic generic axial torsion coupling	70
2.9.17	Cubic viscoelastic generic	71
2.9.18	Double linear viscoelastic	71
2.9.19	Turbulent viscoelastic	71
2.9.20	Linear viscoelastic bistop	72
2.9.21	GRAALL damper	72
2.9.22	shock absorber	72
2.9.23	symbolic elastic	73
2.9.24	symbolic viscous	74
2.9.25	symbolic viscoelastic	74
2.9.26	ann elastic	74
2.9.27	ann viscoelastic	75
2.9.28	nlsf viscoelastic	75
2.9.29	nlp viscoelastic	77
2.9.30	array	78
2.9.31	bistop	79
2.9.32	invariant angular	79
2.10	Hydraulic fluid	79
2.10.1	Incompressible	80
2.10.2	Linearly compressible	80
2.10.3	Linearly compressible, with thermal dependency	80
2.10.4	Super (linearly compressible, with thermal dependency)	80
2.10.5	Exponential compressible fluid, with saturation	81
2.11	Authentication Methods	81
2.11.1	Note on security and confidentiality	82
2.11.2	No authentication	82
2.11.3	Password	82
2.11.4	PAM (Pluggable Authentication Modules)	82
2.11.5	SASL (Simple Authentication and Security Layer)	83
2.12	Miscellaneous	83
3	Data	86
3.1	Problem	86
4	Problems	87
4.1	Initial-Value Problem	87
4.1.1	General Data	87
4.1.2	Derivatives Data	99
4.1.3	Dummy Steps Data	99
4.1.4	Output Data	100
4.1.5	Real-Time Execution	101
4.2	Other Problems	102

5	Control Data	103
5.1	Assembly-Related Cards	103
5.1.1	Skip Initial Joint Assembly	103
5.1.2	Use	103
5.1.3	Initial Stiffness	104
5.1.4	Omega Rotates	104
5.1.5	Tolerance	104
5.1.6	Max Iterations	104
5.2	General-Purpose Cards	104
5.2.1	Title	104
5.2.2	Print	104
5.2.3	Make Restart File	105
5.2.4	Select Timeout	105
5.2.5	Default Output	106
5.2.6	Output Frequency	107
5.2.7	Output Meter	107
5.2.8	Output Precision	107
5.2.9	Output Results	107
5.2.10	Default Orientation	108
5.2.11	Default Aerodynamic Output	108
5.2.12	Default Beam Output	109
5.2.13	Default Scale	109
5.2.14	Model	110
5.2.15	Rigid Body Kinematics	110
5.2.16	Loadable path	111
5.3	Model Counter Cards	111
5.3.1	Nodes	111
5.3.2	Drivers	111
5.3.3	Elements	111
6	Nodes	113
6.1	Structural Node	114
6.1.1	Static Node	114
6.1.2	Dynamic Node	114
6.1.3	Modal Node	114
6.1.4	Syntax	114
6.1.5	Dummy Node	115
6.2	Electric Node	119
6.3	Abstract Node	119
6.4	Hydraulic Node	119
6.5	Parameter Node	120
6.6	Miscellaneous	120
6.6.1	Output	120
7	Drivers	122
7.1	File Drivers	122
7.1.1	Fixed Step	122
7.1.2	Variable Step	124
7.1.3	Socket	124
7.1.4	RTAI Mailbox	125

7.1.5	Stream	125
8	Elements	128
8.1	Aerodynamic Elements	129
8.1.1	Aerodynamic Body/Aerodynamic Beam2/3	129
8.1.2	Aeromodal Element	134
8.1.3	Aircraft Instruments	135
8.1.4	Air Properties	136
8.1.5	Generic Aerodynamic Force	140
8.1.6	Induced velocity	143
8.1.7	Rotor	146
8.2	Automatic structural	146
8.3	Beam Elements	147
8.3.1	Beam Section Constitutive Law	148
8.3.2	Three-node beam element	158
8.3.3	Two-node beam element	162
8.3.4	Output	164
8.3.5	Notes	164
8.4	Body	164
8.5	Bulk Elements	166
8.5.1	Stiffness spring	167
8.6	Couple	167
8.7	Electric Elements	167
8.7.1	Accelerometer	167
8.7.2	Displacement	168
8.7.3	Motor	168
8.7.4	Discrete control	169
8.8	Force	173
8.8.1	Output	174
8.8.2	Abstract force	174
8.8.3	Abstract reaction force	175
8.8.4	Structural force	175
8.8.5	Structural internal force	177
8.8.6	Structural couple	178
8.8.7	Structural internal couple	179
8.8.8	Modal	179
8.8.9	External forces	181
8.8.10	External Structural	183
8.8.11	External structural mapping	186
8.8.12	External modal	190
8.8.13	External modal mapping	193
8.9	Genel Element	195
8.9.1	General Purpose Elements	195
8.9.2	Special Rotorcraft GENEL Elements	200
8.10	Gravity Element	202
8.11	Hydraulic Element	202
8.11.1	Actuator	203
8.11.2	Minor Loss	203
8.11.3	Three Way Minor Loss	203
8.11.4	Control Valve	204

8.11.5	Dynamic Control Valve	204
8.11.6	Pressure Flow Control Valve	205
8.11.7	Pressure Valve	205
8.11.8	Flow Valve	205
8.11.9	Orifice	206
8.11.10	Accumulator	206
8.11.11	Tank	206
8.11.12	Pipe	206
8.11.13	Dynamic Pipe	207
8.11.14	Imposed Pressure	207
8.11.15	Imposed Flow	207
8.12	Joint Element	207
8.12.1	Angular acceleration	208
8.12.2	Angular velocity	208
8.12.3	Axial rotation	209
8.12.4	Beam slider	210
8.12.5	Brake	210
8.12.6	Cardano hinge	211
8.12.7	Cardano pin	212
8.12.8	Cardano rotation	212
8.12.9	Clamp	212
8.12.10	Coincidence	213
8.12.11	Deformable Axial Joint	213
8.12.12	Deformable displacement hinge	214
8.12.13	Deformable displacement joint	214
8.12.14	Deformable Hinge	216
8.12.15	Deformable joint	217
8.12.16	Distance	219
8.12.17	Distance with offset	220
8.12.18	Drive displacement	220
8.12.19	Drive displacement pin	221
8.12.20	Drive hinge	221
8.12.21	Gimbal hinge	222
8.12.22	Gimbal rotation	222
8.12.23	Imposed displacement	223
8.12.24	Imposed displacement pin	224
8.12.25	In line	224
8.12.26	In plane	225
8.12.27	Invariant deformable displacement joint	225
8.12.28	Invariant deformable hinge	225
8.12.29	Kinematic	226
8.12.30	Linear acceleration	226
8.12.31	Linear velocity	226
8.12.32	Modal	227
8.12.33	Plane displacement	230
8.12.34	Plane displacement pin	230
8.12.35	Plane hinge	231
8.12.36	Plane pin	231
8.12.37	Prismatic	231

8.12.38	Revolute hinge	231
8.12.39	Revolute pin	233
8.12.40	Revolute rotation	234
8.12.41	Rod	234
8.12.42	Rod with offset	236
8.12.43	Spherical hinge	236
8.12.44	Spherical pin	237
8.12.45	Total equation	237
8.12.46	Total joint	237
8.12.47	Total pin joint	241
8.12.48	Total reaction	243
8.12.49	Universal hinge	243
8.12.50	Universal pin	243
8.12.51	Universal rotation	243
8.12.52	Viscous body	243
8.12.53	Lower Pairs	244
8.13	Joint Regularization	244
8.13.1	Tikhonov	244
8.14	Plate Elements	245
8.14.1	Shell4	245
8.15	User-Defined Elements	246
8.15.1	Loadable Element	246
8.15.2	User-Defined Element	247
8.15.3	General Discussion on Run-Time Loadable Modules	248
8.16	Output Elements	249
8.16.1	Stream output	249
8.16.2	RTAI output	252
8.16.3	Stream motion output	252
8.17	Miscellaneous	253
8.17.1	Bind	253
8.17.2	Driven Element	255
8.17.3	Inertia	258
8.17.4	Output	258
A	NetCDF Output Format	259
A.1	NetCDF Output	259
A.1.1	Base Level	260
A.1.2	Run Level	260
A.1.3	Node Level	260
A.1.4	Element Level	261
A.2	Accessing the Database	263
A.2.1	Octave	263
B	Modal Element Data	264
B.1	FEM File Format	264
B.1.1	Example: Dynamic Model	267
B.1.2	Example: Static Model	268
B.2	Code Aster Procedure	270
B.3	NASTRAN Procedure	271
B.4	Procedures for Other FEA Software	272

C	Results Visualization	273
C.1	EasyAnim	273
C.2	Output Manipulation	274
C.2.1	Output in a Relative Reference Frame	274
D	Log File Format	276
D.1	Generic Format	276
D.2	Model Description Entries	276
D.2.1	aero0	276
D.2.2	aero2	276
D.2.3	aero3	277
D.2.4	axial rotation	277
D.2.5	beam2	277
D.2.6	beam3	277
D.2.7	brake	278
D.2.8	clamp	278
D.2.9	deformable joints	278
D.2.10	distance	278
D.2.11	gimbal rotation	279
D.2.12	inline	279
D.2.13	inplane	279
D.2.14	prismatic	279
D.2.15	relative frame structural node	279
D.2.16	revolute hinge	280
D.2.17	revolute rotation	280
D.2.18	rod	280
D.2.19	spherical hinge	280
D.2.20	spherical pin	280
D.2.21	structural forces	281
D.2.22	structural node	281
D.2.23	universal hinge	281
D.2.24	universal pin	281
D.2.25	universal rotation	282
D.3	Analysis Description Entries	282
D.3.1	inertia	282
D.3.2	output frequency	282
D.3.3	struct node dofs	283
D.3.4	struct node momentum dofs	283
D.3.5	struct node labels	283
E	Frequently Asked Questions	284
E.1	Input	284
E.1.1	What is the exact syntax of element X ?	284
E.1.2	What element should I use to model Y ?	284
E.2	Output	285
E.2.1	How can I reduce the amount of output?	285
E.3	Execution Debugging	286
E.3.1	How can I find out why the iterative solution of a nonlinear problem does not converge?	286

List of Figures

2.1	Predefined constants in math parser	19
2.2	Construction of 3×3 orientation matrix from two non-parallel vectors.	27
8.1	Airfoil geometry	132
8.2	Constitutive coefficients grouping (S: shear, A: axial)	149
8.3	Beam section	155
8.4	Geometry of the three-node beam element.	158
8.5	Geometry of the two-node beam element.	162
8.6	NASTRAN CONM2 card	165
8.7	Discrete control layout.	170
8.8	Shell: definitions	246

List of Tables

2.1	Built-in mathematical operators in math parser (from higher to lower precedence)	20
2.2	Built-in mathematical functions in math parser	21
2.3	Built-in types in math parser	21
2.4	Predefined variables and constants in math parser	22
2.5	Functions in the <code>model</code> namespace	23
2.6	Drive callers	42
2.7	Constitutive laws dimensionality	65
4.1	Linear solvers properties	98
4.2	Linear solvers memory usage	98
4.3	Linear solvers pivot handling	98

Introduction

This document describes the format of the input data for MBDyn, the Multibody Dynamics analysis suite. It should be considered a syntax more than a format, since the rules of the input allow a lot of freedom in the file format.

As the title states, this manual describes the syntax of the input. The reader should not expect to *learn* how to model systems from this text; it is rather a reference manual where the user can find the exact and detailed description of the syntax of some entity, where it is assumed that the user already knows that entity exactly does what is required to model that system or solve that problem.

To get on the right track, one may have a look at the *Frequently Asked Questions* (Appendix E), a steadily growing chapter that collects the most common issues arising from users. They might help pointing to the right statement for a given purpose, but in any case nothing can replace the judgment of the user about what entity exactly fits some need. The only document that may help users in growing this type of judgment is the *tutorials* book, available from the website. For more specific questions, one should try the

`mbdyn-users@mbdyn.org`

mailing list (browse the archives, just in case, then post a question; posting is restricted to subscribers; subscription is free).

Conventions on the Notation

Throughout the manual, a (sometimes naïve) Backus-Naur-like syntax description is used. Extensions are made, such as to put optional arguments in square brackets ‘[]’, mutually exclusive arguments in curly brackets ‘{}’, separated by the operator ‘|’ (the logical “OR”). Non-terminal symbols are enclosed in angle brackets ‘<>’, while terminal symbols are written in normal types. The association of a non-terminal with terminal or non-terminal symbols is performed by means of the operator ‘:=’. When required, the type of a non-terminal symbol is enforced in a “C” programming language casting style, by prepending a (usually self-explanatory) type enclosed in plain brackets ‘()’.

Remarks

The input of the program MBDyn is subject to continuous changes at a fast pace, since the code is under development. This documentation attempts to give some insight into the logics that drove its implementation.

While most of the software resulted from a careful design, some portions are “as they are” only because they were made in a hurry, or because no better way was at hand at the moment they were made. The input format partially reflects the development of the software. Whenever changes in the input format and in the syntax of existing parts are required, an effort will be attempted to make it as much backward

compatible as possible, or at least reliable diagnostics and error checking will be put in place that warns for possible erroneous inputs related to a change in the input format.

At present the input of MBDyn already warns for many possible errors and makes as many cross-checks as possible on the consistency of the data. Nearly every warning and error issue shows the line of the input file where the error occurred.

This document will be subjected to changes, so be sure you always have a release that is aligned with the version of the code you're running.

For any question or problem, to fix typos, bugs, for comments and suggestions, please contact the users' mailing list

`mbdyn-users@mbdyn.org`

Please note that you need to subscribe to be allowed to post to the list. Posting is public, so please only post information not subjected to distribution restrictions.

Mailing list subscription information is available on the web site

`http://www.aero.polimi.it/mbdyn/`

As an alternative, you can directly contact the Developers' Team:

Pierangelo Masarati,
MBDyn Development Team
Dipartimento di Ingegneria Aerospaziale
Politecnico di Milano
via La Masa 34, 20156 Milano, Italy
Fax: +39 02 2399 8334
E-mail: `mbdyn@aero.polimi.it`
Web: `http://www.aero.polimi.it/mbdyn/`

This manual is also available online at
`http://www.aero.polimi.it/masarati/MBDyn-input/manual/index.html`

Make sure you use the manual that refers to the version of MBDyn that you are using; from MBDyn 1.2 on, the input manual is marked with the same version number of the package.

The website `http://www.mbdyn.org/` may advertise different versions of the manual, including those related to past releases and a HEAD or -devel version, which refers to current development; this can be useful to check whether the current development may address problems you are having. Moreover, since the updating of the manual is not as prompt as required to keep pace with code releases, the -devel manual may describe features that are already in the latest release but not yet documented in that release's manual.

Chapter 1

Overview

This chapter gives a global overview of the structure of the input file.

1.1 Execution

MBDyn is a command-line tool. This means that in order to execute it, one needs to start the executable from a terminal, and, as a minimum, pass it some indications about the problem it must analyze.

Usually, this is done by preparing one or more input files that describe the model and the analysis that needs to be performed.

1.1.1 Passing the Input

The input file can be passed either using a specific switch (`-f`, preferred), or by listing multiple input file names on the command line. In the latter case, the analyses are executed in sequence. The format of the input file is the topic of most of this manual. If no input file is defined, MBDyn as a last resort tries to read from the standard input. As a consequence, the following commands are equivalent:

```
# assuming file "input" exists in the current working directory and is readable...
$ mbdyn -f input
$ mbdyn input
$ mbdyn < input
```

They are listed in order of preference.

If the input file name consists of a path, either absolute or relative, the file can reside anywhere in the file system. In this case, MBDyn changes its current working directory (see `chdir(2)` for details) to the directory where the input file resides. As a consequence, the path of any included file is seen as relative to the directory where the file containing the `include` directive is located (see Section 2.3.5 for details).

1.1.2 Passing the Output File Name

The output is stored in a family of files whose common name is passed through another switch (`-o`). Those files are formed by adding an extension specific to their contents to the output file name. If none is defined, the output file name defaults to `MBDyn`.

If the argument of the `-o` switch consists of a path, either absolute or relative, and a file name, the full path up to the file name excluded must exist and be write-accessible. The output files are stored in

the corresponding directory. Otherwise they are stored in the current working directory of the process at the time of its execution.

If the argument of the `-o` switch consists of exactly a directory, either absolute or relative, without any file name, the name of the input file is used, but the output files are placed in the directory indicated in the argument of the `-o` switch.

For example:

```
$ mbdyn -f input                # output in "input.<extension>"
$ mbdyn -f input -o output      # output in "output.<extension>"
$ mbdyn -f input -o folder/output # output in "folder/output.<extension>"
$ mbdyn -f input -o folder/     # output in "folder/input.<extension>"
$ mbdyn < input                # output in "MBDyn.<extension>"
```

1.2 Input File Structure

The input is divided in blocks. This is a consequence of the fact that almost every module of the code needs data and each module is responsible for its data input. So it is natural to make each module read and interpret its data starting from the input phase.

Every statement (or ‘card’) follows the syntax:

```
<card> ::= <description> [ : <arglist> ] ;
```

`arglist` is a(n optional) list of arguments, that is driven by the `description` that identifies the card. The keyword can be made of multiple words separated by spaces or tabs; the extra spaces¹ are skipped, and the match is case insensitive. The arguments are usually separated by commas². A semicolon ends the card.

Many cards allow extra arguments that may assume default values if not supplied by the user. Usually these arguments are at the end of the card and must follow some rules. A check for the existence of extra args is made, and they are usually read in a fixed sequence. Optional args are usually prefixed by a keyword.

When structured arguments like matrices, vectors, or drive callers are expected, they are parsed by dedicated functions. Anyway, the structured data always follows the rules of the general data. Few exceptions are present, and will be eliminated soon. Every data block is surrounded by the control statements `begin` and `end`:

```
begin : <block_name> ;
      # block content
end   : <block_name> ;
```

The block sequence is:

```
begin : data ;
      # select a problem
      problem : <problem_name> ;
end   : data ;

begin : <problem_name> ;
      # problem-specific data
```

¹ Anything that makes the C function `isspace()` return `true`

² A few exceptions require a colon to separate blocks of arguments; wherever it is feasible, those exceptions will be eliminated in future versions. Those cards will be marked as deprecated.

```

end : <problem_name> ;

begin : control data ;
    # model control data
end : control data ;

begin : nodes ;
    # nodes data
end : nodes ;

# note: optional, if requested in control data
begin : drivers ;
    # special drivers data
end : drivers ;

begin : elements ;
    # elements data
end : elements ;

```

The `drivers` block is optional.

The Schur solver allows an extra block after that of the elements:

```

# note: optional; only allowed when Schur solver is defined in problem_name
begin : parallel ;
    # parallel data
end : parallel ;

```

where data specific to the partitioning and the connectivity of the Schur solver is provided. This is yet undocumented, and will likely be described in a future chapter.

Chapter 2 describes the basic and aggregate data structures that concur in building up the cards. Chapter 3 describes problem selection options. Chapter 4 describes the cards dedicated to the specification of the parameters of the simulation and of the selected integration scheme. Chapter 5 describes the model control data. Chapter 6 describes the cards related to the input of nodes. Nodes come first because they are the elementary blocks of the connectivity graph of the model, so historically they had to be defined before any element could be introduced. Chapter 7 describes the cards related to special drivers. Chapter 8 describes the cards related to the input of elements.

1.3 Output

The program outputs data to a set of files for each simulation. The contents of each file is related to the file extension, which is appended to the input file if no output file name is explicitly supplied.

If no input file is explicitly provided as well, and thus input is directly read from `stdin`, the output file name defaults to ‘MBDyn’. Otherwise, unless the output file name is explicitly set, the name of the input file is used.

The contents of the output files are described within the description of the items (nodes or elements) that generate them. Only a general information file, with extension `.out`, is described here. The file contains general information about the simulation; it is not formatted.

The file contains occasional informational messages, prefixed by a hash mark (`#`). These messages should be intended as comments about the current status of the simulation. At some point, after initialization completes, the comment

```
# Step Time TStep NIter ResErr SolErr SolConv
```

appears, which illustrates the contents of the lines that will be written for each time step. The fields indicate:

1. **Step**: the time step number;
2. **Time**: the time at that step;
3. **TStep**: the time step at that step ($\text{Time}_k - \text{Time}_{k-1}$);
4. **NIter**: the number of iterations required to converge;
5. **ResErr**: the error on the residual after convergence (0 if not computed);
6. **SolErr**: the error on the solution after convergence (0 if not computed);
7. **SolConv**: a boolean that indicates whether convergence was determined by the error criterion on the residual or on the solution (0 for residual, 1 for solution).

There is also a supplementary file, with `.log` extension, that may contain extra (logging) information. Its content, although very experimental and subjected to changes, is documented in Appendix D.

Experimental support for output using the NetCDF database format is available for selected items. See Appendix A for details.

Chapter 2

General

This chapter describes how data structures are read and how they participate, as building blocks, to the definition of specific cards. Consistency across the software and the input file has been a driving principle in designing the input of MBDyn. As such, the very same elementary data structures are present in very different contexts.

2.1 Types

2.1.1 Keywords

In MBDyn's input a lot of so-called "keywords" appear literally in the input file. They are case-insensitive and may have an arbitrary amount of space in between. For example, `'control data'`, `'Control Data'`, and `'controldata'` are equivalent.

However, most of the keywords are context-dependent, so they are not illustrated altogether in a dedicated section, but rather presented in relation to the contexts they may appear in.

2.1.2 Strings

Literal strings are not used very much in MBDyn. However, they may be used in quite a few significant places, for example to indicate file names, so few details on their syntax are provided.

Strings are typically delimited by double quotes (""). When a string is expected, the parser looks for the opening double quotes and eats up all the white space before it. Then all characters are read as they are until the closing double quotes are encountered. The escape character is the backslash ('\'); it is used:

- to escape the escape character itself ('\');
- to escape any double quotes that are part of the string, and would otherwise be interpreted as the closure of the string;
- to break a string on multiple lines by placing it before the newline character ('\n'); in this case, the escape character and the newline are eaten up;
- to allow the use of non-printing characters, represented in the form '`\<hexpair>`', so that the hexadecimal representation of the non-printing char is converted into its integer equivalent in the 0–255 range.

Character sets different from ASCII are not supported.

2.1.3 Numeric Values

Every time a numeric value is expected, the result of evaluating a mathematical expression can be used, including variable declaration and assignment (variable names and values are kept in memory throughout the input phase and the simulation) and simple math functions. Named variables and non-named constants are strongly typed; four types are currently available: **bool**, **integer**, **real**, and **string**.

Operations are strongly typed and perform implicit cast when allowed. For instance `1+2.5` returns a **real** whose value is `3.5`, since one of the two addenda is **real**, while `1/3` returns `0` because both values are implicit integers and thus the integer division is used.

An empty field, delimited by a valid separator (a comma or a semicolon, depending on whether other arguments are expected or not) returns the (program supplied) default value for that field, if supplied by the caller, otherwise the parser automatically defaults to zero.

Multiple expressions can be used, provided they are enclosed in plain brackets and are separated by semicolons; the result of the last expression will be used as the expected numeric value, but all the expressions (which may have persistent effects, like variable declarations and assignments) will be evaluated.

Example.

```
1.                                     # == 1.
(real r = 2.*pi; integer i = 3; sin(i*r*Time+.87)) # == sin(6*pi*Time+.87)
```

The latter sequence evaluates to $\sin(6\pi t + .87)$, depending on the current value of the pre-defined variable **Time**. The constant **pi** is always defined as π with machine precision (the macro `M_PI` in the standard C header file `math.h`). Other constants are pre-defined, as illustrated in Table 2.4.

When MBDyn is invoked with the `-H` option,

```
mbdyn -H
```

it prints the predefined variables table. A typical output is shown in Figure 2.1.3. Real values are stored with the maximum precision allowed by the underlying **real** type (by default, double precision, 64 bit).

The variable **Time** is declared, defined and initialized¹ from the beginning of the control data section, and during the solution phase it is assigned the value of the current time.

Table 2.1 shows the supported mathematical operators, while Table 2.2 shows the built-in mathematical functions. The supported types are listed in Table 2.3. Table 2.4 lists the predefined variables; notice that they're treated exactly as user-defined variables, so they can be reassigned.

2.1.4 Namespaces

The math parser uses the notion of *namespace* to separate functions. The functions listed in Table 2.2 are implicitly defined in the **default** namespace, i.e. they should be referenced by writing

```
default::sqrt(2.)
```

Other namespaces may be defined by the user, and loaded by means of the **module load** card described in Section 2.3.6.

The namespace that refers to the current model is loaded by default. Its name is **model**, and contains the functions defined in Table 2.5. With the exception of the **drive** and **sf** functions, all functions in the **model** namespace with the suffix `_prev` (e.g. `position_prev`) use the configuration at the previous

¹ Note: a variable is *declared* when its name enters the namespace; it is *defined* when it can be referenced; it is *initialized* when it is first assigned a value.

```
user@host:~> mbdyn -H

MBDyn - MultiBody Dynamics 1.X-Devel
compiled on Jul 25 2010 at 21:36:59

Copyright 1996-2012 (C) Paolo Mantegazza and Pierangelo Masarati,
Dipartimento di Ingegneria Aerospaziale <http://www.aero.polimi.it/>
Politecnico di Milano <http://www.polimi.it/>

MBDyn is free software, covered by the GNU General Public License,
and you are welcome to change it and/or distribute copies of it
under certain conditions. Use 'mbdyn --license' to see the conditions.
There is absolutely no warranty for MBDyn. Use "mbdyn --warranty"
for details.

default symbol table:
  const integer RAND_MAX = 2147483647
  const real deg2rad = 0.0174533
  const real e = 2.71828
  const real ft2m = 0.3048
  const real in2m = 0.0254
  const real in2mm = 25.4
  const real kg2lb = 2.20462
  const real kg2slug = 0.0685218
  const real lb2kg = 0.453592
  const real m2ft = 3.28084
  const real m2in = 39.3701
  const real mm2in = 0.0393701
  const real pi = 3.14159
  const real rad2deg = 57.2958
  const real slug2kg = 14.5939

MBDyn terminated normally
```

Figure 2.1: Predefined constants in math parser

Table 2.1: Built-in mathematical operators in math parser (from higher to lower precedence)

Operator	Type	Description
^	Binary, right	Power
+	Unary, left	Plus sign
-	Unary, left	Minus sign
*	Binary, left	Multiplication
/	Binary, left	Division
+	Binary, left	Addition
-	Binary, left	Subtraction
>	Binary, left	Greater than
>=	Binary, left	Greater than or equal to
==	Binary, left	Equal to
<=	Binary, left	Less than or equal to
<	Binary, left	Less than
!=	Binary, left	Not equal
!	Unary, right	NOT
&&	Binary, left	AND
	Binary, left	OR
~	Binary, left	XOR (exclusive OR)
=	Binary, right	Assignment

Note: “left” and “right” refer to the associativity of the operators.

Note: the `string` type only supports the binary “+” operator.

time step. They may be useful, for example, when parsing hints (see Section 2.5.37), since `hint` parsing occurs after prediction, but often needs to refer to the configuration at the previous, just completed time step.

2.1.5 Plugin Variables

Plugin variables, or plugins, as the name states, are pluggable extensions to the math parser that allow to register mechanisms to bind a variable to some means to dynamically generate its value. As a consequence, any time a variable declared as part of a plugin is used, its value is dynamically evaluated executing the related code.

There are built-in plugins that allow to link variables to the value of degrees of freedom and to special parameters computed by nodes and elements.

The syntax of a plugin consists in specially defining a variable so that it gets registered to the appropriate code:

```
set : OSQBR <plugin> , <var> [ , <arglist> ] CSQBR ;
<arglist> ::= <arg> [ , ... ]
```

where OSQBR and CSQBR stand for ‘open’ and ‘close’ square brackets, respectively, which should not be confused with those that indicate optional parameters.

In the following, the built-in plugins² are illustrated. For details on nodes, elements, their properties and the information they can expose, see the related sections.

²Built-in and plug-in may sound contradictory, and in fact they are. However, consider that the mathematical parser is essentially independent of MBDyn (in fact, it has been turned into a command-line calculator in `c1(1)`), so, from this point of view, built-in plug-ins are used to plug MBDyn information into an independent, lower-level piece of code.

Table 2.2: Built-in mathematical functions in math parser

Name	Returns	Arg[s]	Description
<code>asin</code>	Real	Real	Arc sine
<code>acos</code>	Real	Real	Arc cosine
<code>atan</code>	Real	Real	Arc tangent
<code>actan</code>	Real	Real	Arc co-tangent
<code>atan2</code>	Real	Real, Real	(Robust) arc tangent
<code>actan2</code>	Real	Real, Real	(Robust) arc co-tangent
<code>cos</code>	Real	Real	Cosine
<code>sin</code>	Real	Real	Sine
<code>tan</code>	Real	Real	Tangent
<code>ctan</code>	Real	Real	Co-tangent
<code>cosh</code>	Real	Real	Hyperbolic cosine
<code>sinh</code>	Real	Real	Hyperbolic sine
<code>tanh</code>	Real	Real	Hyperbolic tangent
<code>ctanh</code>	Real	Real	Hyperbolic co-tangent
<code>acosh</code>	Real	Real	Hyperbolic arc cosine
<code>asinh</code>	Real	Real	Hyperbolic arc sine
<code>atanh</code>	Real	Real	Hyperbolic arc tangent
<code>actanh</code>	Real	Real	Hyperbolic arc co-tangent
<code>exp</code>	Real	Real	Exponential
<code>log</code>	Real	Real	Natural logarithm
<code>log10</code>	Real	Real	Base 10 logarithm
<code>sqrt</code>	Real	Real	Square root
<code>abs</code>	Real	Real	Absolute value
<code>sign</code>	Real	Real	Sign
<code>copysign</code>	Real	Real, Real	First arg with sign of second
<code>floor</code>	Integer	Real	Closest integer from below
<code>ceil</code>	Integer	Real	Closest integer from above
<code>round</code>	Integer	Real	Closest integer
<code>rand</code>	Integer	Void	random integer [0, RAND_MAX]
<code>random</code>	Real	Void	random real [-1.0, 1.0]
<code>seed</code>	Void	Integer	Seeds the random number generator
<code>step</code>	Real	Real	Step function
<code>ramp</code>	Real	Real	Ramp function
<code>sramp</code>	Real	Real, Real	Saturated ramp function
<code>par</code>	Real	Real	Parabolic function
<code>print</code>	Void	Real	Prints a value to standard output

Table 2.3: Built-in types in math parser

Name	Description
<code>bool</code>	Boolean number (promoted to <code>integer</code> , <code>real</code> , or <code>string</code> (0 or 1), whenever required)
<code>integer</code>	Integer number (promoted to <code>real</code> , or <code>string</code> , whenever required)
<code>real</code>	Real number (promoted to <code>string</code> whenever required)
<code>string</code>	Text string

Table 2.4: Predefined variables and constants in math parser

Name	Type	Value
<code>Time</code>	Real	Current simulation time
<code>TimeStep</code>	Real	Current simulation time step
<code>Step</code>	Integer	Current simulation step
<code>Var</code>	Real	Set by <code>dof</code> , <code>node</code> , or <code>element</code> drive callers with degree of freedom value, node or element private data value, respectively
<code>e</code>	Real	Neper's number
<code>pi</code>	Real	π
<code>RAND_MAX</code>	Integer	Maximum random integer
<code>in2m</code>	Real	Inch to meter ratio (0.0254)
<code>m2in</code>	Real	Meter to inch ratio (1.0/0.0254)
<code>in2mm</code>	Real	Inch to meter ratio (25.4)
<code>mm2in</code>	Real	Meter to inch ratio (1.0/25.4)
<code>ft2m</code>	Real	Foot to meter ratio (0.3048)
<code>m2ft</code>	Real	Meter to foot ratio (1.0/0.3048)
<code>lb2kg</code>	Real	Pound to kilogram ratio (0.45359237)
<code>kg2lb</code>	Real	Kilogram to pound ratio (1.0/0.45359237)
<code>deg2rad</code>	Real	Degree to radian ratio ($\pi/180$)
<code>rad2deg</code>	Real	Radian to degree ratio ($180/\pi$)
<code>slug2kg</code>	Real	Slug to kilogram ratio (14.5939)
<code>kg2slug</code>	Real	Kilogram to slug ratio (1.0/14.5939)

Table 2.5: Functions in the `model` namespace

Name	Ret.	Arg[s]	Description
<code>position</code>	Real	Integer	norm of structural node position
<code>position2</code>	Real	Integer	square norm of structural node position
<code>xposition</code>	Real	Integer	X component of structural node position
<code>yposition</code>	Real	Integer	Y component of structural node position
<code>zposition</code>	Real	Integer	Z component of structural node position
<code>distance</code>	Real	Integer, Integer	distance between structural nodes
<code>distance2</code>	Real	Integer, Integer	square distance between structural nodes
<code>xdistance</code>	Real	Integer, Integer	X component of distance between structural nodes
<code>ydistance</code>	Real	Integer, Integer	Y component of distance between structural nodes
<code>zdistance</code>	Real	Integer, Integer	Z component of distance between structural nodes
<code>xunitvec</code>	Real	Integer, Integer	X component of unit vector between structural nodes
<code>yunitvec</code>	Real	Integer, Integer	Y component of unit vector between structural nodes
<code>zunitvec</code>	Real	Integer, Integer	Z component of unit vector between structural nodes
<code>angle</code>	Real	Integer	norm of rotation vector of structural node
<code>xangle</code>	Real	Integer	X component of structural node rotation vector
<code>yangle</code>	Real	Integer	Y component of structural node rotation vector
<code>zangle</code>	Real	Integer	Z component of structural node rotation vector
<code>anglerel</code>	Real	Integer, Integer	angle between structural nodes (norm of rotation vector)
<code>xanglerel</code>	Real	Integer, Integer	X component of rotation vector between structural nodes
<code>yanglerel</code>	Real	Integer, Integer	Y component of rotation vector between structural nodes
<code>zanglerel</code>	Real	Integer, Integer	Z component of rotation vector between structural nodes
<code>velocity</code>	Real	Integer	norm of structural node velocity
<code>velocity2</code>	Real	Integer	square norm of structural node velocity
<code>xvelocity</code>	Real	Integer	X component of structural node velocity
<code>yvelocity</code>	Real	Integer	Y component of structural node velocity
<code>zvelocity</code>	Real	Integer	Z component of structural node velocity
<code>vrel</code>	Real	Integer, Integer	norm of relative velocity between structural nodes
<code>vrel2</code>	Real	Integer, Integer	square norm of relative velocity between structural nodes
<code>xvrel</code>	Real	Integer, Integer	X component of relative velocity between structural nodes
<code>yvrel</code>	Real	Integer, Integer	Y component of relative velocity between structural nodes
<code>zvrel</code>	Real	Integer, Integer	Z component of relative velocity between structural nodes
<code>angvel</code>	Real	Integer	norm of structural node angular velocity
<code>angvel</code>	Real	Integer	square norm of structural node angular velocity
<code>xangvel</code>	Real	Integer	X component of structural node angular velocity
<code>yangvel</code>	Real	Integer	Y component of structural node angular velocity
<code>zangvel</code>	Real	Integer	Z component of structural node angular velocity
<code>angvrel</code>	Real	Integer, Integer	norm of relative angular velocity between structural nodes
<code>angvrel2</code>	Real	Integer, Integer	square norm of rel. angular velocity between structural nodes
<code>xangvrel</code>	Real	Integer, Integer	X component of rel. angular velocity between structural nodes
<code>yangvrel</code>	Real	Integer, Integer	Y component of rel. angular velocity between structural nodes
<code>zangvrel</code>	Real	Integer, Integer	Z component of rel. angular velocity between structural nodes
<code>current</code>	Real	String	allows to retrieve data specific to the current context (essentially undocumented)
<code>drive</code>	Real	Integer, Real	evaluates the <code>drive caller</code> indicated by its label (the first argument) for the input specified as second argument
<code>sf::<name></code>	Real	Real, (opt)Integer	evaluates the <code>scalar function name</code> for the input specified as argument (value, optional derivative order)
<code>node::<type></code>	Real	Integer, String	evaluates a private datum of a node (label, string)
<code>element::<type></code>	Real	(opt)Integer, String	evaluates a private datum of an element (label, string)

Dof plugin

Note: this plugin is now obsolete by the `node` plugin, which allows to access more significant information about nodes, including the value of the degrees of freedom.

The `dof` plugin allows to link a variable to a specific degree of freedom of a node, or to its derivative, whenever defined. The syntax is

```
<plugin> ::= dof

<arglist> ::= <label> , <type>
           [ , <index> ] ,
           { algebraic | differential }
           [ , prev= <prev> ]

<prev> ::= { true | false }
```

where `label` is the label of the node, `type` is the node type, `index` is the index of the degree of freedom that is requested, if the node has more than one, while `algebraic` and `differential` refer to the value of the dof and of its derivative, respectively. The derivative can be requested only for those degrees of freedom that have one. When `prev` is `true`, the value at the previous time step is used.

Example. The variable `VARNAME` takes the value of the derivative of `abstract` node `NODENAME`

```
set: integer NODENAME = 1000;
# the node must exist
set: [dof,VARNAME,NODENAME,abstract,differential];
```

Node plugin

The `node` plugin allows to link a variable to any data a node can expose, including the values of its degrees of freedom. The syntax is

```
<plugin> ::= node
<arglist> ::= <label> , <type>
           [ , { [ index= ] <index> | string= <name> } ]
```

where `label` is the label of the node, `type` is the node type, and the mutually exclusive `index` and `name` represent the index of the datum that is requested, `name` being a user-friendly representation of the actual index. The `index` form is the default; however, the `name` form is recommended. Note that `name` should be enclosed in double quotes, although not strictly required. In fact, some node types allow data names that include square brackets. In those cases, double-quote enclosing is needed to avoid parsing errors, since closing square brackets indicate the end of the plugin variable specification.

Example. This implements a variable `VARNAME` with exactly the same value of the one defined in the example of the `dof` plugin, where the variable takes the value of the derivative of the `abstract` node `NODENAME`

```
set: integer NODENAME = 1000;
# the node must exist
set: [node,VARNAME,NODENAME,abstract,string="xP"];
```

Element plugin

The `element` plugin allows to link a variable to any data an element can expose, including the values of its degrees of freedom. The syntax is

```
<plugin> ::= element
<arglist> ::= <label> , <type>
           [ , { [ index= ] <index> | string= <name> } ]
```

where `label` is the label of the element, `type` is the element type, and the mutually exclusive `index` and `name` represent the index of the datum that is requested, `name` being a user-friendly representation of the actual index. The `index` form is the default; however, the `name` form is recommended. Note that `name` should be enclosed in double quotes, although not strictly required. In fact, some element types allow names that include square brackets. In those cases, double-quote enclosing is needed to avoid parsing errors, since closing square brackets indicate the end of the plugin variable specification.

Example. The variable `VARNAME` takes the value of the z component of the reaction force of `joint` `ELEMNAME`

```
set: integer ELEMNAME = 1000;
# the joint must exist
set: [element,VARNAME,ELEMNAME,joint,string="Fz"];
```

2.2 Higher-Order Math Structures

Every time a higher-order mathematical structure is expected, it can be preceded by a keyword that influences how the structure is read. All of the available structures support the keyword `null` which causes the structure to be initialized with zeros. When a non-null value is input, it can be followed by the keyword `scale` with a scale factor; the scale factor can be any mathematical expression. This is useful to rescale structure values by reassigning the value of the scale factor. The main data structures are:

2.2.1 3×1 vectors

1. general case: a sequence of 3 reals, comma-separated.
2. null vector: keyword `null`; the vector is initialized with zeros.

As an example, all the following lines define an empty 3×1 vector:

```
default
null
0.,0.,0.
,,
```

the first case is correct if no default was actually available for that specific vector, thus falling back to three zeros. The following rescales an arbitrary vector

```
cos(pi/3.),0.,sin(pi/3.), scale, 100.
```

2.2.2 6×1 vectors

1. general case: a sequence of 6 reals, comma-separated.
2. null vector: keyword **null**; the vector is initialized with zeros.

2.2.3 3×3 matrices

1. general case: a sequence of 9 reals, comma-separated, which represent the row-oriented coefficients $a_{11}, a_{12}, \dots, a_{32}, a_{33}$. *Note: the 9 coefficients can be preceded by the keyword **matr** for consistency with other entities; its use is recommended whenever an ambiguity is possible.*

$$\mathbf{m} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \quad (2.1)$$

2. symmetric matrix: keyword **sym**, followed by a sequence of 6 reals, comma-separated, that represents the upper triangle, row-oriented coefficients of a symmetric matrix, e.g. $a_{11}, \dots, a_{13}, a_{22}, a_{23}, a_{33}$.

$$\mathbf{m} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{12} & a_{22} & a_{23} \\ a_{13} & a_{23} & a_{33} \end{bmatrix} \quad (2.2)$$

3. skew symmetric matrix: keyword **skew**, followed by a sequence of 3 reals, comma-separated, that are the components of the vector \mathbf{v} that generates a skew symmetric matrix $\mathbf{m} = \mathbf{v} \times$:

$$\mathbf{m} = \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{bmatrix} \quad (2.3)$$

4. diagonal matrix: keyword **diag**, followed by a sequence of 3 reals, comma-separated, that represent the diagonal coefficients of a diagonal matrix, namely a_{11}, a_{22}, a_{33}

$$\mathbf{m} = \begin{bmatrix} a_{11} & 0 & 0 \\ 0 & a_{22} & 0 \\ 0 & 0 & a_{33} \end{bmatrix} \quad (2.4)$$

5. identity matrix: keyword **eye**; the matrix is initialized as the identity matrix, that is a null matrix except for the diagonal coefficients that are 1.
6. null matrix: keyword **null**; the matrix is initialized with zeros.

For example, the identity matrix can be defined as:

```
matr, 1.,0.,0., 0.,1.,0., 0.,0.,1.
1.,0.,0., 0.,1.,0., 0.,0.,1.      # ‘‘matr’’ omitted
sym, 1.,0.,0., 1.,0., 1.          # upper triangular part
diag, 1.,1.,1.                    # diagonal
eye                                # a la matlab
```

Although not required, for better readability it is recommended to format the above data as




Figure 2.2: Construction of 3×3 orientation matrix from two non-parallel vectors.

```

matr,
  1.,0.,0.,
  0.,1.,0.,
  0.,0.,1.
1.,0.,0.,
0.,1.,0.,
0.,0.,1.          # ‘‘matr’’ omitted
sym,
  1.,0.,0.,
  1.,0.,
  1.              # upper triangular part
diag, 1.,1.,1.   # diagonal
eye              # a la matlab

```

2.2.4 3×3 orientation matrices

1. general case: two vectors that define an orthonormal reference system, each of them preceded by its index in the final orientation matrix. The procedure that builds the orientation matrix from the two vectors is illustrated in Figure 2.2. The first vector, vector 1 in subfigure (a), is normalized and assumed to represent the desired direction. The second vector, vector 2' in subfigure (b), simply concurs in the definition of the plane the vector that is not given is normal to. The third vector, vector 3 in subfigure (c), results from the cross-product of the first two vectors, after normalization. The actual direction of the second vector, vector 2 in subfigure (d), results from the cross product of the third and the first vector. Examples:

```

1, 1.,0.,0., 2, 0.,1.,0.
1, (real ALPHA = pi/6.; cos(ALPHA)), sin(ALPHA), 0.,
3, 0.,0.,1.

```

The first example represents the identity matrix, i.e. no rotation occurs with respect to the global reference frame: direction 1 in the local frame is parallel to $1.,0.,0.$, which represents direction

1 in the global frame, while direction 2 in the local frame is parallel to $0., 1., 0.$, which represents direction 2 in the global frame.

The second example describes a rotation of $\pi/6$ radian about global direction 3: direction 1 in the local frame results from composing `cos(pi/6.)` in global direction 1 and `sin(pi/6.)` in global direction 2, while direction 3 in the local frame remains parallel to $0., 0., 1.$, which represents direction 3 in the global frame.

2. a variant of the above, which may be useful when only one direction really matters, is illustrated in the example below:

```
1, 1., 0., 0., 2, guess
```

The keyword `guess` tells the parser to generate a random vector that is orthogonal to the given one, which is used as the direction indicated by the index (2 in the example). The vector is computed based on a very simple algorithm: it contains

- 1.0 corresponding to the index with smallest module, v_1 (min);
- $-v_1$ (min) / v_1 (max) corresponding to the index with the largest module, v_1 (max);
- 0.0 on the remaining index.

3. identity matrix: keyword `eye`; the identity matrix, which means there is no rotation with respect to the global reference frame.
4. a complete orientation matrix: keyword `matr` followed by the nine, row-oriented, coefficients, namely $r_{11}, r_{12}, \dots, r_{33}$. *Note: no orthogonality check is performed; be sure an orthogonal matrix, within the desired tolerance, is input.*
5. Euler angles: keyword `euler`, followed by the three values, as output by structural nodes. The keywords `euler123`, `euler313` and `euler321` allow to use orientations in the sequence specified by the keyword.

```
{ euler | euler123 | euler313 | euler321 } ,
  [ degrees , ]
  <angle_1> , <angle_2> , <angle_3>
```

Note: the output in the `.mov` file is in degrees, while in input MBDyn requires angles in radians, unless the values are preceded by the keyword `degrees`.

Note: the definition of the three angles that are used by the code to express orientations may vary between versions. Currently, Bryant-Cardano angles are used in place of Euler angles; see the note related to the output of the structural nodes in Section 6.1.5, and the Technical Manual. The code will remain consistent, i.e. the same angle definition will be used for input and output, but models over versions may become incompatible, so this syntax should really be used only as a means to quickly reproduce in the input an orientation as resulting from a previous analysis.

6. Orientation vector: keyword `vector`, followed by the three values, as output by structural nodes. The resulting matrix is computed as

$$\mathbf{R} = \exp(\mathbf{v} \times) \quad (2.5)$$

2.2.5 6×6 matrices

1. general case: a sequence of 36 reals, comma-separated, that represent the row-oriented coefficients $a_{11}, a_{12}, \dots, a_{65}, a_{66}$.
2. ANBA format: keyword **anba**, followed by 36 reals, comma-separated, that represent the coefficients of the beam stiffness matrix as generated by the beam section analysis code ANBA, namely the following transformation is performed:
 - axis x , in the section plane in ANBA notation, becomes axis 2 in MBDyn notation;
 - axis y , in the section plane in ANBA notation, becomes axis 3 in MBDyn notation;
 - axis z , the beam axis in ANBA notation, becomes axis 1 in MBDyn notation;

Note: this format is mainly intended for backwards compatibility with older versions of that beam section analysis software, which used a different numbering convention for the reference frame that is local to the beam section.
3. symmetric matrix: keyword **sym**, followed by a sequence of 21 reals, comma-separated, that represents the upper triangle, row-oriented coefficients of a symmetric matrix, e.g. $a_{11}, \dots, a_{16}, a_{22}, \dots, a_{26}, \dots, a_{66}$.
4. diagonal matrix: keyword **diag**, followed by a sequence of 6 reals, comma-separated, that represent the diagonal coefficients of a diagonal matrix.
5. identity matrix: keyword **eye**; the matrix is initialized as the identity matrix, that is a null matrix except for the diagonal coefficients that are 1.
6. null matrix: keyword **null**; the matrix is initialized with zeros.

2.2.6 $6 \times N$ matrices

1. general case: a sequence of $6 \times N$ reals, comma-separated, that represent the row-oriented coefficients $a_{11}, a_{12}, \dots, a_{6(N-1)}, a_{6N}$.
2. ANBA format: keyword **anba**, followed by $6 \times N$ reals, comma-separated, that represent the coefficients of the beam stiffness matrix as generated by the code ANBA, namely the following transformation is performed:
 - axis x , in the section plane in ANBA notation, becomes axis 2 in MBDyn notation;
 - axis y , in the section plane in ANBA notation, becomes axis 3 in MBDyn notation;
 - axis z , the beam axis in ANBA notation, becomes axis 1 in MBDyn notation;
3. null matrix: keyword **null**; the matrix is initialized with zeros.

2.3 Input Related Cards

(Almost) everywhere in the input file the statement cards defined in the following can be used. They are handled directly by the parsing object, and merely act as an indirect reference to entities that are not explicitly enumerated. They are:

2.3.1 Constitutive Law

```
<card> ::= constitutive law : <label> ,  
  [ name , " <name> " , ]  
  <dim> , (ConstitutiveLaw<dim>D) <constitutive_law> ;
```

Constitutive laws are grouped by their dimensionality `dim`, which (up to now) can be any of 1, 3 and 6. The `constitutive_law` is parsed according to the rules described in Section 2.9, and can be referenced later when it needs to be used.

2.3.2 C81 Data

This keyword allows to define and read the `c81 data` airfoil tables that are used by aerodynamic elements.

```
<card> ::= c81 data : <label> [ , name , " <name> " ]  
  " <filename> "  
  [ , tolerance , <tolerance> ]  
  [ , { free format | nrel | fc511 } ]  
  [ , flip ]  
  [ , echo , " <output_filename> " [ , free format ] ] ;
```

Data is read from file `filename` according to the format specified in the following.

The optional keyword `tolerance` allows to specify the tolerance that is used to determine the boundaries of the linear portion of the lift curve slope.

The traditional format is used unless a format is specified using any of the keywords `free format`, `nrel` or `fc511` (the latter is intentionally not documented). The format is automatically inferred by reading the heading line of the input file.

The optional keyword `flip` instructs MBDyn to “flip” the data, i.e. treat the airfoil data as if the sign of the angle of attack changed along with that of the lift and moment coefficients, but not that of the drag coefficient. This corresponds to considering the airfoil “upside down”.

The optional keyword `echo` allows to specify the name of the file that will be generated with the data just read, for cross-checking purposes. If the following optional keyword is `free format`, or if data was read in free format, the echo will also be in free format. Otherwise, it will be in the traditional format.

Traditional Format

The file is in textual form; the traditional format (the default) is:

- first line: "%30s%2d%2d%2d%2d%2d" where the first 30 chars are a title string, currently ignored by MBDyn, followed by 6 consecutive two-digit integers that indicate:
 - the number ML of *Mach* points for C_l ;
 - the number NL of angle of attack points for C_l ;
 - the number MD of *Mach* points for C_d ;
 - the number ND of angle of attack points for C_d ;
 - the number MM of *Mach* points for C_m ;
 - the number NM of angle of attack points for C_m .

The example in `var/naca0012.c81` contains:

```
1.....01.....01.....0MLNLMDNDMMNM; not part of format  
PROFILO NACA 0012          11391165 947
```

- the format of each following line is up to 10 fields of 7 chars each; records longer than 10 fields are broken on multiple lines, with the first field filled with blanks;
- a block containing the C_l data, made of:
 - a record with the first field blank, followed by the ML *Mach* values for the C_l ;
 - NL records containing the angle of attack in the first field, followed by ML values of C_l for each *Mach* number; angles of attack wrap around 360 deg, starting from -180 deg.

The example provided in file `var/naca0012.c81` contains 11 *Mach* points and 39 angle of attack records for the C_l of the NACA0012 airfoil at high (unspecified) *Reynolds*:

```

.....7.....7.....7.....7.....7.....7.....7.....7.....7.....7.....7.....7.....7; not part of format
      0.    .20   .30   .40   .50   .60   .70   .75   .80
      .90   1.
-180.  0.    0.    0.    0.    0.    0.    0.    0.    0.
      0.    0.
-172.5 .78   .78   .78   .78   .78   .78   .78   .78   .78
      .78   .78
...

```

- a block containing the C_d data, same as for C_l , with MD *Mach* points and ND angle of attack records;
- a block containing the C_m data, same as for C_l , with MM *Mach* points and NM angle of attack records.

Free Format

Finally, to allow higher precision whenever available, a native format, based on `c81`, called **free format**, is available. It basically consists in the `c81` format without continuation lines and with arbitrary precision, with fields separated by blanks. The header is made of an arbitrary string, terminated by a semicolon ‘;’, followed by the six numbers that define the dimensionality of the expected data.

Example.

```

# FREE FORMAT
this is the header; 2 8 2 2 2 2
      0.0  0.9
-180.0  0.0  0.0
-170.0  1.0  0.9
 -90.0  0.0  0.0
-10.0 -1.0 -0.9
  10.0  1.0  0.9
  90.0  0.0  0.0
 170.0 -1.0 -0.9
 180.0  0.0  0.0
      0.0  0.9
-180.0  0.1  0.1
  180.0  0.1  0.1
      0.0  0.9
-180.0  0.0  0.0
  180.0  0.0  0.0

```

NREL

The `nrel` format is used to provide airfoil data for wind turbines. The format is very compact, but does not allow to account for Mach dependence, and requires the same angle of attack resolution for all three coefficients. See AeroDyn's User Guide [1] for a description of the format. Please note that MBDyn only supports one airfoil per file, and only considers static coefficients from line 15 on.

To be completed.

Alternative Format

An alternative format, required by some projects, can be used by supplying the optional switch `fc511`; it is intentionally not documented.

Format Conversion

Traditional format files can be automatically converted in free format using the `c81test(1)` utility:

```
$ c81test -d var/naca0012_free_format.c81 var/naca0012.c81
```

generates a free format version of the NACA 0012 airfoil data provided in the distribution.

Miscellaneous

A simple program that allows to read and plot the C81 tables is available at <http://homepage.mac.com/jhuwaldt/java/Applications/TableReader/TableReader.html> (thanks to Marco Fossati for pointing it out).

A simple utility that parses the C81 file and computes the aerodynamic coefficients for a given pair of angle of attack and Mach number is `c81test(1)`, available in the MBDyn utils suite. This routine uses exactly the same code internally used by MBDyn, so it should be considered a check of the code rather than a real tool.

2.3.3 Drive Caller

The `drive caller` directive

```
<card> ::= drive caller : <label> ,  
        [ name , " <name> " , ]  
        (DriveCaller) <drive_caller> ;
```

allows to define a `drive caller` (see Section 2.5) that can be subsequently reused. It is useful essentially in two cases:

- a) to define a `drive` that will be used many times throughout a model;
- b) to define a `drive` that needs to be used in a later defined part of a model, in order to make it parametric.

2.3.4 Hydraulic fluid

The `hydraulic fluid` directive:

```
<card> ::= hydraulic fluid : <label> ,  
        <fluid_type> , <fluid_properties> ;
```

allows to define a hydraulic fluid to be later used in hydraulic elements, see Section 8.11. The fluid is identified by the `label`. The `fluid_types`, with the related `fluid_properties`, are described in 2.10

2.3.5 Include

The `include` directive

```
<card> ::= include : " <file_name> " ;
```

allows to include the contents of the file `file_name`, which must be a valid filename for the operating system in use. The file name must be enclosed in double quotes ("). The full (absolute or relative) path must be given if the included file is not in the directory of the including one.

There is no check for recursive `include` directives, so **the user must take care of recursion.**

The `include` directive forces the parser to scan the included file `file_name` before continuing with the including one. This is very useful if, for instance, a big model can be made of many small models that are meaningful by themselves.

It can be used to replicate parts of the model, by simply using parametric labels for nodes, elements, reference systems, and setting a bias value before multiple-including the same bulk data file. Examples of this usage are given in the tutorials (<http://www.aero.polimi.it/mbdyn/documentation/tutorials/>).

2.3.6 Module Load

The `module load` directive:

```
<card> ::= module load : " <file_name> "  
      [ , <module_arglist> ] ;
```

```
<module_arglist> ::= <arg> [ , ... ]
```

where `file_name` is the name of a runtime loadable object, causes the object to be opened, and a function `module_init()`, with prototype

```
extern "C" int  
module_init(const char *module_name,  
            void *data_manager, void* mbdyn_parser);
```

to be executed.

The function is assumed to perform the operations required to initialize the module, eventually taking advantage of the parsing and of the data manager; see the technical manual for details.

The function is also expected to take care of the optional `module_arglist` arguments; in detail, module developers are encouraged to support `help` as the first optional argument. The presence of this argument should result in printing to standard output as much information as possible about the use of the module.

The typical use consists in registering some methods for later use. A clear example is given in

```
modules/module-wheel2/module-wheel2.cc
```

where the `module_init()` function registers the `user defined` element called `wheel2` in the set of user-defined element handlers. The module can be retrieved later using the syntax described in Section 8.15. Note, however, that the `module_init()` function may be used for any purpose. A typical use consists in registering any kind of handlers for subsequent use.

Run-time module loading is taken care of by GNU's `libltdl`. Modules are compiled using `libtool` for portability purposes. The resulting modules take the name `libmodule-<name>.la`.

Modules are installed by default in the directory $\${prefix}/libexec$. When loaded using only the module name, the default directory is searched. The run-time loading path can be modified by the `loadable path` statement described in Section 5.2.16.

Although `libltdl` is supposed to be portable on a wide variety of platforms (this is what it is designed for, all in all), run-time loading within MBDyn is mainly tested using Linux. Users are encouraged to report problems they might encounter, especially when building modules for different platforms, as this would help making MBDyn more portable.

2.3.7 Print symbol table

The `print symbol table` directive:

```
<card> ::= print symbol table ;
```

allows to print to standard output the contents of the parser's symbol table at any stage of the input phase. This may be useful for model debugging purposes.

2.3.8 Reference

The `reference` directive:

```
<card> ::= reference : <unique_label> ,
          (Vec3) <absolute_position> ,
          (OrientationMatrix) <absolute_orientation_matrix> ,
          (Vec3) <absolute_velocity> ,
          (Vec3) <absolute_angular_velocity> ;
```

A `reference` system is declared and defined. It must be given a unique identifier, scanned by the math parser (which means that any regular expression is allowed, and the result is rounded up to the nearest unsigned integer). The entries `absolute_*` are parsed by routines that compute absolute (i.e. referring to the global frame) entities starting from a given entity in a given reference frame. These routines are very general, and make intense use of the `reference` entries themselves, which means that a reference can be recursively defined by means of previously defined `reference` entries.

Alternative form:

```
<card> ::= reference : <unique_label> ,
          [ reference , { global | <ref_label> } , ]
          <d> , <theta> , <a> , <alpha> ;
```

where:

- **d**: offset d along previous z axis to the common normal (the current x axis);
- **theta**: the angle θ about previous z from previous to current x axis;
- **a**: length a of the common normal; for revolute joints, the radius of the previous z ;
- **alpha**: the angle α about the common normal, from previous to new z axis.

The new reference data is

$$\mathbf{x}_{\text{new}} = \mathbf{x}_{\text{old}} + \mathbf{R}_{\text{old}} (d\mathbf{e}_3 + \exp((\theta\mathbf{e}_3) \times) \cdot (a\mathbf{e}_1)) \quad (2.6a)$$

$$\mathbf{R}_{\text{new}} = \mathbf{R}_{\text{old}} \exp((\theta\mathbf{e}_3) \times) \exp((\alpha\mathbf{e}_1) \times) \quad (2.6b)$$

$$\mathbf{v}_{\text{new}} = \mathbf{v}_{\text{old}} + \boldsymbol{\omega}_{\text{old}} \times \mathbf{R}_{\text{old}} (d\mathbf{e}_3 + \exp((\theta\mathbf{e}_3) \times) \cdot (a\mathbf{e}_1)) \quad (2.6c)$$

$$\boldsymbol{\omega}_{\text{new}} = \boldsymbol{\omega}_{\text{old}} \quad (2.6d)$$

Use of Reference Frames

Every time an absolute or a relative geometric or physical entity is required, it is processed by a set of routines that allow the entity to be expressed in the desired reference frame. The following cases are considered:

- relative position (physical)
- absolute position (physical)
- relative orientation matrix (physical)
- absolute orientation matrix (physical)
- relative velocity (physical)
- absolute velocity (physical)
- relative angular velocity (physical)
- absolute angular velocity (physical)
- relative arbitrary vector (geometric)
- absolute arbitrary vector (geometric)

The caller is responsible for the final interpretation of the input. The caller always supplies the routines a default reference structure the input must be referred to. So, depending on the caller, the entry can be in the following forms:

1. **<entity>**:
the data supplied in **entity** is intended in the default reference frame
2. **reference** , **<reference_type>** , **<entity>**:
the data is in **reference_type** reference frame, where

$$\langle \text{reference_type} \rangle ::= \{ \text{global} \mid \text{node} \mid \text{local} \}$$
3. **reference** , **<reference_label>** , **<entity>**:
the data is in **reference_label** reference frame. This reference frame must be already defined.

In some cases, significantly in case of joints that connect two nodes a and b , special reference types are allowed when reading specific entities related to the second node. These reference types compute the value of the entity with respect to the reference frame associated to the first node, a :

1. **other position**:
a relative position $\tilde{\mathbf{p}}$ is intended in the other node's reference frame, with respect to the relative position $\tilde{\mathbf{f}}_a$ already specified for the other node,

$$\tilde{\mathbf{f}}_b = \mathbf{R}_b^T \left(\mathbf{x}_a + \mathbf{R}_a \left(\tilde{\mathbf{f}}_a + \tilde{\mathbf{p}} \right) - \mathbf{x}_b \right), \quad (2.7)$$

which is the solution of equation

$$\mathbf{x}_a + \mathbf{R}_a \left(\tilde{\mathbf{f}}_a + \tilde{\mathbf{p}} \right) = \mathbf{x}_b + \mathbf{R}_b \tilde{\mathbf{f}}_b; \quad (2.8)$$

2. **other orientation:**

a relative orientation $\tilde{\mathbf{R}}$ is intended in the other node's reference frame, with respect to the relative orientation $\tilde{\mathbf{R}}_{ha}$ already specified for the other node,

$$\tilde{\mathbf{R}}_{hb} = \mathbf{R}_b^T \mathbf{R}_a \tilde{\mathbf{R}}_{ha} \tilde{\mathbf{R}}, \quad (2.9)$$

which is the solution of equation

$$\mathbf{R}_a \tilde{\mathbf{R}}_{ha} \tilde{\mathbf{R}} = \mathbf{R}_b \tilde{\mathbf{R}}_{hb}; \quad (2.10)$$

3. **other node:**

a relative position $\tilde{\mathbf{p}}$ or a relative orientation $\tilde{\mathbf{R}}$ are intended in the other node's reference frame,

$$\tilde{\mathbf{f}}_b = \mathbf{R}_b^T (\mathbf{x}_a + \mathbf{R}_a \tilde{\mathbf{p}} - \mathbf{x}_b) \quad (2.11a)$$

$$\tilde{\mathbf{R}}_{hb} = \mathbf{R}_b^T \mathbf{R}_a \tilde{\mathbf{R}}, \quad (2.11b)$$

which are the solutions of equations

$$\mathbf{x}_a + \mathbf{R}_a \tilde{\mathbf{p}} = \mathbf{x}_b + \mathbf{R}_b \tilde{\mathbf{f}}_b \quad (2.12a)$$

$$\mathbf{R}_a \tilde{\mathbf{R}} = \mathbf{R}_b \tilde{\mathbf{R}}_{hb}. \quad (2.12b)$$

Example.

- absolute position:

```

null
reference, global, null
reference, 8, 1., sin(.3*pi), log(3.)

```

- relative orientation matrix (e.g. as required by many constraints and thus referred to a node):

```

eye
reference, node, eye
reference, 8,
    3, 0., 1., 0.,
    1, .5, sqrt(3)/2., 0.

```

Notes:

- the global reference frame has position $\{0, 0, 0\}$, orientation matrix **eye**, velocity $\{0, 0, 0\}$ and angular velocity $\{0, 0, 0\}$.
- if the caller is not related to a node, the reference type **node** should not be defined. In this case it is considered equivalent to **local**.
- when processing a velocity or an angular velocity, the resulting value always accounts for the velocity and angular velocity of the frame the entry is referred to.

As an example, if a node is defined in a reference frame \mathbf{R}_R that has non-null angular velocity $\boldsymbol{\Omega}_R$, and the position $\mathbf{x}_{\text{input}}$ of the node is not coincident with the origin \mathbf{X}_R of the reference frame it

is attached to, its global velocity and angular velocity result as the composition of the input values and of those of the reference frame:

$$\begin{aligned}\mathbf{w} &= \mathbf{R}_R \boldsymbol{\omega}_{\text{input}} + \boldsymbol{\Omega}_R \\ \mathbf{v} &= \mathbf{R}_R \mathbf{v}_{\text{input}} + \mathbf{V}_R + \boldsymbol{\Omega}_R \times (\mathbf{R}_R \mathbf{x}_{\text{input}})\end{aligned}$$

This, for instance, eases the input of all the parts of a complex system that is moving as a rigid body, by defining a reference frame with the proper initial velocities, and then referring all the entities, e.g. the nodes, to that frame, with null local velocity.

Recalling the declaration and the definition of reference frames, a simple reference frame definition, with all the entries referring by default to the global system, would be:

```
reference: 1000,
        null,
        eye,
        null,
        null;
```

which represents a redefinition of the global system. A more verbose, and self-explanatory definition would be:

```
reference: 1000,
        reference, global, null,
        reference, global, eye,
        reference, global, null,
        reference, global, null;
```

the reference frame one is referring to must be repeated for all the entries since they must be allowed to refer to whatever frame is preferred by the user. A fancier definition would be:

```
reference: Rotating_structure,
        reference, Fixed_structure, null,
        reference, Spindle_1,
            1, 0.,0.,1.,
            3, 0.,1.,0.,
        reference, Fixed_structure, null,
        reference, Spindle_1, 0.,0.,Omega_1;
```

Output

The reference frames are used only during the input phase, where they help referring entities either absolute or relative to other entities depending on their internal representation during the analysis. As such, reference frames cannot be “used” or “visualized” neither directly nor indirectly at any time during the analysis or by interpreting the output, because they do not “evolve” nor are attached to any state-dependent entity. To allow their debugging, however, they can be output in the global reference frame according to the representation of structural nodes, as described in Section 6.1.5, by using the **default output** directive with the value **reference frames**, as detailed in Section 5.2.5.

2.3.9 Remark

The `remark` directive:

```
<card> ::= remark : " <remark_string> "  
          [ , <math_expression> [ , ... ] ] ;
```

This directive simply prints to stdout the string `remark_string` and optionally evaluates and prints any subsequent expression `math_expression` according to the `set` directive. It is used to allow rough input debugging, where the file name and line is logged, followed by a message and by the evaluation of any expression.

Example. A file “remarks”, containing only the statements

```
remark: "square root of 2", sqrt(2);  
set: (  
    real EA = 1e6; # N, axial stiffness  
    real GA = 1e6; # N, shear stiffness  
    real EJ = 1e3; # Nm^2, bending stiffness  
    real GJ = 1e3; # Nm^2, torsional stiffness  
0);  
remark: "Stiffness properties", EA, GA, EJ, GJ;
```

results in

```
user@host:~>$ mbdyn -f remarks
```

```
MBDyn - Multi-Body Dynamics 1.2.5  
compiled on Nov 18 2005 at 15:15:39
```

```
Copyright 1996-2012 (C) Paolo Mantegazza and Pierangelo Masarati,  
Dipartimento di Ingegneria Aerospaziale <http://www.aero.polimi.it/>  
Politecnico di Milano <http://www.polimi.it/>
```

```
MBDyn is free software, covered by the GNU General Public License,  
and you are welcome to change it and/or distribute copies of it  
under certain conditions. Use 'mbdyn --license' to see the conditions.  
There is absolutely no warranty for MBDyn. Use "mbdyn --warranty"  
for details.
```

```
reading from file "remarks"  
line 1, file <remarks>: square root of 2, 1.41421  
line 8, file <remarks>: Stiffness properties, 1e+06, 1e+06, 1000, 1000  
MBDyn terminated normally  
user@host:~>$
```

2.3.10 Set

The `set` directive:

```
<card> ::= set : <math_expression> ;
```

This directive simply invokes the math parser to evaluate the expression `math_expression` and then discards the result.

It can be used to declare new variables, or to set the values of existing ones. This feature is very useful, since it allows to build parametric models, and to reuse generic model components.

Example.

```
set: integer LABEL = 1000;    # a label
set: real E = 210e+9;        # Pa, steel's elastic module
set: real A = 1e-4;         # m^2, a beam's cross section
set: real K = E*A;          # N, a beam's axial stiffness
```

2.3.11 Setenv

The `setenv` directive:

```
<card> ::= setenv : [ overwrite , ] " <varname> [ = <value> ] " ;
```

This directive sets the environment variable `varname` to `value`, if given; otherwise, the variable is unset. If the keyword `overwrite` is set, the variable is overwritten, if already set.

```
# set FILE to "test", if it does not exist
setenv: "FILE=test";
# set FILE to "test", even if it exists
setenv: override, "FILE=test";
# unset FILE
setenv: "FILE";
# set FILE to the empty string, if it does not exist
setenv: "FILE=";
```

See `setenv(3)` and `unsetenv(3)` man pages for details.

2.3.12 Template Drive Caller

The `template drive caller` directive

```
<card> ::= template drive caller : <label> , { " <dim_name> " | (int) <dim> } ,
      (TplDriveCaller<Entity>) <tpl_drive_caller> ;
```

```
<dim_name> ::= { 1 | 3 | 6 | 3x3 | 6x6 }
```

```
<dim> ::= { 1 | 3 | 6 }
```

allows to define a `template drive caller` (see Section 2.5.38) that can be subsequently reused. The type is defined according to

Entity	dim_name	dim
double real	1	1
Vec3	3	3
Vec6	6	6
Mat3x3	3x3	
Mat6x6	6x6	

Template drive callers are useful essentially in two cases:

- a) to define a **template drive** that will be used many times throughout a model;
- b) to define a **template drive** that needs to be used in part of a model that is defined later, in order to make it parametric.

2.4 Node Degrees of Freedom

A node in MBDyn is an entity that owns public degrees of freedom and instantiates the corresponding public equations. It can lend them to other entities, called elements, to let them write contributions to public equations, possibly depending on the value of the public degrees of freedom.

Usually elements access nodal degrees of freedom through well-defined interfaces, at a high level. But in a few cases, nodal degrees of freedom must be accessed at a very low level, with the bare knowledge of the node label, the node type, the internal number of the degree of freedom, and the order of that degree of freedom (algebraic or differential, if any). The data that allows an entity to track a nodal degree of freedom is called `NodeDof`; it is read according to

```
<node_dof> ::= <node_label> ,  
            <node_type>  
            [ , <dof_number> ]  
            [ , { algebraic | differential } ]
```

The label `node_label` and the type `node_type` of the node are used to track the pointer to the desired node.

If `node_type` refers to a non-scalar node type, the `dof_number` field is required to indicate the requested degree of freedom.

Finally, the order of the degree of freedom is read, if required. It must be one of **algebraic** or **differential**. If the `dof_number` degree of freedom is differential, both orders can be addressed, while in case of a node referring to an algebraic degree of freedom there is no choice, only the **algebraic** order can be addressed and thus this field is not required.

The `dof_number` must be between 1 and the number of degrees of freedom related to the node. Not all numbers might be valid for specific nodes; for example, `dof_number` values 4, 5, and 6 are not valid for a **structural node**, Section 6.1 when the order is **algebraic**.

When the node degree of freedom input syntax is used to address an equation (as in **abstract** force elements, Section 8.8.2, or in **discrete control** elements, Section 8.7.4), the distinction between **algebraic** and **differential** is meaningless, and thus this field is not required.

2.5 Drives and Drive Callers

Every time some entity can be “driven”, i.e. a value can be expressed as dependent on some “external” input, an object of the class `DriveCaller` is used. The **drive** essentially represents a scalar function, whose value can change over time or, through some more sophisticated means, can depend on the state of the analysis. Usually, the dependence over time is implicitly assumed, unless otherwise specified. For example, the amplitude of the force applied by a **force** element (see Section 8.8) is defined by means of a **drive**; as such, the value of the **drive** is implicitly calculated as a function of the time. However, a **dof drive** (see Section 2.5.7) uses a subordinate **drive** to compute its value based on the value of a degree of freedom of the analysis; as a consequence, the value of the **dof drive** is represented by the value of the subordinate **drive** when evaluated as a function of that specific degree of freedom at the desired time (function of function).

The family of the `DriveCaller` object is very large. The type of the `DriveCaller` is declared as

```

<drive_caller> ::=
  { <drive_caller_type> [ , <arglist> ]
    | reference , <label> }

```

where `arglist`, if any, is a comma-separated list of arguments that depends on `drive_caller_type`. As an exception, a constant `DriveCaller` (that behaves exactly as a numerical constant with little or no overhead depending on the optimizing capability of the compiler) is assumed when a numeric value is used instead of a keyword.

If the alternative format is used, the keyword `reference` must be followed by the label of an already defined, valid drive caller (See Section 2.3.3).

Drive callers are listed in Table 2.6.

2.5.1 Array drive

```

<drive_caller> ::= array ,
  <num_drives> ,
  (DriveCaller) <drive_caller>
  [ , ... ]

```

this is simply a front-end for the linear combination of `num_drives` normal drives; `num_drives` must be at least 1, in which case a simple drive caller is created, otherwise an array of drive callers is created and at every call their value is added to give the final value of the array drive.

2.5.2 Closest next drive

```

<drive_caller> ::= closest next ,
  <initial_time> ,
  { forever | <final_time> } ,
  (DriveCaller) <increment>

```

This drive returns a non-zero value when called for the first time with an argument greater than or equal to the current threshold value, which is computed starting from `initial_time` and incrementing it each time by as many `increment` values as required to pass the current value of `Time`. As soon as a threshold value is exceeded, as many values of `increment` as required to pass the current value of `Time` are added, and the process repeats.

This drive caller is useful within the `output meter` statement (see Section 5.2.7) to obtain nearly equally spaced output when integrating with variable time step.

2.5.3 Const(ant) drive

```

<drive_caller> ::= [ const , ] <const_coef>

```

The keyword `const` can be omitted, thus highlighting the real nature of this driver, that is completely equivalent to a constant, static real value.

2.5.4 Cosine drive

```

<drive_caller> ::= cosine ,
  <initial_time> ,
  <angular_velocity> ,
  <amplitude> ,

```

Table 2.6: Drive callers

Name	Differentiable	Notes
array	✓	depends on subordinate drive callers
closest next	✓	depends on subordinate drive caller
const	✓	
cosine	✓	
cubic	✓	
direct	✓	
dof		depends on subordinate drive caller
double ramp	✓	
double step	✓	
drive	✓	depends on subordinate drive callers
element	✓	depends on subordinate drive caller
exponential	✓	
ginac	✓	symbolic differentiation using GiNaC
file		
fourier series	✓	
frequency sweep	✓	depends on subordinate drive callers
linear	✓	
meter	✓	
mult	✓	depends on subordinate drive callers
node	✓	depends on subordinate drive caller
null	✓	
parabolic	✓	
periodic	✓	depends on subordinate drive caller
piecewise linear	✓	
ramp	✓	
random	✓	
sample and hold		
scalar function	✓	depends on underlying scalar function
sine	✓	
step	✓	
string	✓	
tanh	✓	
time	✓	
timestep	✓	
unit	✓	

```
{ [ - ] <number_of_cycles> | half | one | forever } ,
<initial_value>
```

where `angular_velocity` is $2\pi/T$. This drive actually computes a function of the type

$$f(t) = \text{initial_value} + \text{amplitude} \cdot (1 - \cos(\text{angular_velocity} \cdot (t - \text{initial_time}))).$$

The value of `number_of_cycles` determines the behavior of the drive. If it is positive, `number_of_cycles` oscillations are performed. If it is negative, the oscillations end after `number_of_cycles`-1/2 cycles at the top of the cosine, with null tangent. Special keywords can be used for `number_of_cycles`:

- `forever`: the oscillation never stops;
- `one`: exactly one cycle is performed;
- `half`: exactly half cycle is performed, so the function stops at

$$t = \text{initial_time} + \frac{\pi}{\text{angular_velocity}},$$

with a value of

$$f = 2 \cdot \text{amplitude} + \text{initial_value},$$

starting and ending with continuous derivative. In this case, it may be convenient to indicate the `angular_velocity` as $\pi/\text{duration}$, where `duration` is the time required to complete the half-wave.

TODO: add a plot that exemplifies the three cases.

2.5.5 Cubic drive

```
<drive_caller> ::= cubic ,
  <const_coef> ,
  <linear_coef> ,
  <parabolic_coef> ,
  <cubic_coef>
```

The function

$$\begin{aligned} f(t) = & \text{const_coef} \\ & + \text{linear_coef} \cdot t \\ & + \text{parabolic_coef} \cdot t^2 \\ & + \text{cubic_coef} \cdot t^3. \end{aligned} \tag{2.13}$$

2.5.6 Direct

```
<drive_caller> ::= direct
```

Transparently returns the input value; the `arglist` is empty. It is useful in conjunction with those drive callers that require their output to be fed into another drive caller, like the `dof`, `node` and `element` drive callers, when the output needs to be used as is.

2.5.7 Dof drive

```
<drive_caller> ::= dof ,
  (NodeDof) <driving_dof> ,
  (DriveCaller) <func_drive>
```

a `NodeDof` (see Section 2.4), namely the reference to a degree of freedom of a node, is read. Then a recursive call to a drive data is read. The driver returns the value of the `func_drive drive` using the value of the `NodeDof` as input instead of the time. This can be used as a sort of explicit feedback, to implement fancy springs (where a force is driven through a function by the displacement of the node it is applied to) or an active control system; e.g.:

```
..., dof, 1000, structural, 3, algebraic,
  linear, 0., 1. ...
```

uses the value of the third component (z) of structural node 1000 as is (that is, in a linear expression with null constant coefficient and unit linear coefficient, while

```
..., dof, 1000, abstract, differential,
  string, "2.*exp(-100.*Var)" ...
```

uses the value of the derivative of abstract node 1000 in computing a string expression. Refer to the description of a `NodeDof` (see Section 2.4) entry for further details.

The same effect can be obtained using the `dof` plugin as follows:

```
set: [dof,x,1000,structural,1,algebraic];
# ...
..., string, "2.*exp(-100.*x)" ...
```

which applies a couple whose amplitude is computed by evaluating a `string` drive which depends on variable `x`; this, in turn, is defined as a `dof` plugin, which causes its evaluation in terms of the selected degree of freedom of the node at each invocation.

2.5.8 Double ramp drive

```
<drive_caller> ::= double ramp ,
  <a_slope> ,
  <a_initial_time> ,
  <a_final_time> ,
  <d_slope> ,
  <d_initial_time> ,
  { forever | <d_final_time> } ,
  <initial_value>
```

The function

$$\begin{aligned}
\langle f_a \rangle &::= a_slope \cdot (a_final_time - a_initial_time) \\
\langle f_d \rangle &::= d_slope \cdot (d_final_time - d_initial_time) \\
f(t) &= initial_value
\end{aligned} \tag{2.14}$$

$$+ \begin{cases} 0 & t < a_initial_time \\ a_slope \cdot (t - a_initial_time) & a_initial_time \leq t \leq a_final_time \\ f_a & a_final_time < t < d_initial_time \\ f_a + d_slope \cdot (t - d_initial_time) & d_initial_time \leq t \leq d_final_time \\ f_a + f_d & d_final_time < t \end{cases}$$

2.5.9 Double step drive

```

<drive_caller> ::= double step ,
  <initial_time> ,
  <final_time> ,
  <step_value> ,
  <<initial_value>

```

The function

$$\begin{aligned}
f(t) &= initial_value \\
+ \begin{cases} 0 & t < initial_time \\ step_value & initial_time \leq t \leq final_time \\ 0 & final_time < t \end{cases}
\end{aligned} \tag{2.15}$$

2.5.10 Drive drive

```

<drive_caller> ::= drive ,
  (DriveCaller) <drive_caller1> ,
  (DriveCaller) <drive_caller2>

```

This is simply a “function of function” drive: the output of `drive_caller2` is fed to `drive_caller1` and the result is returned. So the value of `drive_caller2` becomes the input argument to `drive_caller1`. Note that the very same thing occurs, for instance, in the `dof`, `node` and `element` drives, where the value of the `dof` or of the element’s private datum, respectively, are fed into the given drive.

Example.

```

force: 1, abstract,
  1, abstract,
  drive,

```

```

    string, "Var*(Var>0.)",
    sine, 0., 2*pi/.2, 10., forever, 0.;

```

is equivalent to

```

set: real v;
force: 1, abstract,
      1, abstract,
      string, "v=sin(5.*2*pi*Time); 10.*v*(v>0)";

```

2.5.11 Element drive

```

<drive_caller> ::= element ,
  <label> ,
  <type> ,
  [ { string , " <name> " | index , <index> } , ]
  (DriveCaller) <func_drive>

```

a reference to the private data of an element is read. This is made of: the element's `label`, the element's `type` and a specification of which private data is being referred; the `index` can be directly given, prepended by the keyword `index`, or the symbolic name `name` can be used, prepended by the keyword `string`. If that element allows only one private data, the specification can be omitted. Then a recursive call to a drive data is read. The driver returns the value of the `func_drive drive` using the value of the element's private data as input instead of the time. This can be used as a sort of explicit feedback, to implement fancy springs (where a force is driven through a function by the rotation of a joint) or an active control system; e.g.:

```

element, 1000, joint, string, "rz", direct

```

uses the value of the rotation about axis z of a revolute hinge as is (that is, in a linear expression with null constant coefficient and unit linear coefficient, while

```

element, 1000, joint, index, 1
  string, "2.*exp(-100.*Var)"

```

uses the same value, addressed in an alternative manner, in computing a string expression.

The same effect can be obtained using the `element` plugin as follows:

```

set: [element,x,1000,joint,string=rz];
# ...
couple: 1, conservative, 1, 0.,0.,1.,
  string, "2.*exp(-100.*x)";

```

which applies a couple whose amplitude is computed by evaluating a `string` drive which depends on the variable `x`; this, in turn, is defined as an `element` plugin, which causes its evaluation in terms of the element's private data at each invocation.

2.5.12 Exponential drive

```

<drive_caller> ::= exponential ,
  <amplitude_value> ,
  <time_constant_value> ,
  <initial_time> ,
  <initial_value>

```

This drive yields a function that resembles the response of a first-order system to a step input. Its value corresponds to `initial_value` for $t < \text{initial_time}$. For $t \geq \text{initial_time}$, it grows to `initial_value+amplitude_value` exponentially. The growth rate is governed by `time_constant_value`. If `time_constant_value` < 0 , the function diverges.

$$f(t) = \text{initial_value} + \text{amplitude_value} \cdot \left(1 - e^{\left(-\frac{t - \text{initial_time}}{\text{time_constant_value}} \right)} \right) \quad (2.16)$$

Example.

```
..., exponential, 10.0, 0.2, 5.0, 0.0, ...
```

yields a function identical to the response of the system

$$H(s) = \frac{1}{1 + 0.2s} \quad (2.17)$$

to a step input occurring at time $t = 5.0$ with amplitude $A = 10.0$.

2.5.13 File drive

The `DriveCaller` is attached to a file drive object that must be declared and defined in the `drivers` section of the input file (see Section 7). As a consequence, `file` drive callers can only be instantiated after the `drivers` section of the input file.

```
<drive_caller> ::= file ,
  <drive_label>
  [ , <column_number> ]
  [ , amplitude , <amplitude> ]
```

`drive_label` is the label of the `drive` the `DriveCaller` is attached to, while `column_number` is the number of the column the `DriveCaller` refers to (defaults to 1). An additional scaling factor `amplitude` can be used to rescale the drive value (defaults to 1.0).

2.5.14 Fourier series drive

```
<drive_caller> ::= fourier series ,
  <initial_time> ,
  <angular_velocity> ,
  <number_of_terms> ,
  <a_0> ,
  <a_1> , <b_1> ,
  [ ... , ]
  { <number_of_cycles> | one | forever } ,
  <initial_value>
```

This drive corresponds to a Fourier series of fundamental angular velocity ω , truncated after n terms, over a given number of cycles P and starting at a given initial time t_0 as

$$f(t) = \frac{a_0}{2} + \sum_{k=1}^n (a_k \cos(k\omega(t-t_0)) + b_k \sin(k\omega(t-t_0)))$$

for

$$t_0 \leq t \leq t_0 + P \frac{2\pi}{\omega}$$

The value of $f(t)$, if defined, is added to `initial_value`. The value of `number_of_cycles` determines the behavior of the drive. If it is positive, `number_of_cycles` oscillations are performed. Special keywords can be used for `number_of_cycles`:

- **forever**: the oscillation never stops;
- **one**: exactly one cycle is performed;

the `number_of_terms` must be at least 1.

2.5.15 Frequency sweep drive

```
<drive_caller> ::= frequency sweep ,
  <initial_time> ,
  (DriveCaller) <angular_velocity_drive> ,
  (DriveCaller) <amplitude_drive> ,
  <initial_value> ,
  { forever | <final_time> } ,
  <final_value>
```

this drive recursively calls two other drives that supply the angular velocity and the amplitude of the oscillation. Any drive can be used.

$$\begin{aligned}
 \omega(t) &= \text{angular_velocity_drive} \\
 A(t) &= \text{amplitude_drive} \\
 t_i &= \text{initial_time} \\
 t_f &= \text{final_time} \text{ (forever: } +\infty) \\
 f_i &= \text{initial_value} \\
 f_f &= \text{final_value} \\
 f(t) &= \begin{cases} f_i & t < t_i \\ f_i + A(t) \sin(\omega(t)(t - t_i)) & t_i \leq t \leq t_f \\ f_f & t_f < t \end{cases} \quad (2.18)
 \end{aligned}$$

2.5.16 GiNaC

```
<drive_caller> ::= ginac ,
  [ symbol , " <symbol> " , ]
  " <expression> "
```

The function `expression` is evaluated and differentiated, if needed, as a function of the variable passed as the optional `symbol`. If none is passed, `Var` is used. Currently there is no way to share the variables of the math parser used by the `string` drive caller. Not to be confused with the `string` drive caller.

Example.

```
..., ginac, "10*log(1 + Var)" ...
..., ginac, symbol, "x", "10*log(1 + x)" ...
```

2.5.17 Linear drive

```
<drive_caller> ::= linear ,  
    <const_coef> ,  
    <slope_coef>
```

The function

$$f(t) = \text{const_coef} + \text{slope_coef} \cdot t \quad (2.19)$$

2.5.18 Meter drive

The **meter** drive has value zero except for every **steps** steps, where it assumes unit value.

```
<drive_caller> ::= meter ,  
    <initial_time> ,  
    { forever | <final_time> }  
    [ , steps , <steps_between_spikes> ]
```

the first optional entry, preceded by the keyword **steps**, sets the number of steps between spikes.

Example.

```
..., meter, 0., forever, steps, 10, ...
```

defines a drive whose value is zero except at time step 0. and every 10 time steps from that, where it assumes unit value.

2.5.19 Mult drive

The **mult** drive multiplies the value of two subordinate drives.

```
<drive_caller> ::= mult ,  
    (DriveCaller) <drive_1> ,  
    (DriveCaller) <drive_2>
```

2.5.20 Node drive

```
<drive_caller> ::= node ,  
    <label> ,  
    <type> ,  
    [ { string , " <name> " | index , <index> } , ]  
    (DriveCaller) <func_drive>
```

a reference to the private data of a node is read. This is made of: the node's **label**, the node's **type** and a specification of which private data is being referred; the **index** can be directly given, prepended by the keyword **index**, or the symbolic name **name** can be used, prepended by the keyword **string**. If that node allows only one private data, the specification can be omitted. Then a recursive call to a drive data is read. The driver returns the value of the **func_drive** drive using the value of the node's private data as input instead of the time. This can be used as a sort of explicit feedback, to implement fancy springs (where a force is driven through a function by the rotation of a joint) or an active control system; e.g.:

```
..., node, 1000, structural, string, "X[3]",
  linear, 0., 1. ...
```

uses the value of the displacement in direction z (3) of a **structural node**, Section 6.1.5, as is (that is, in a linear expression with null constant coefficient and unit linear coefficient, see the **linear** drive, Section 2.5.17), while

```
..., node, 1000, structural, index, 3,
  string, "2.*exp(-100.*Var)" ...
```

uses the same value, addressed in an alternative manner, in computing a string expression.

2.5.21 Null drive

```
<drive_caller> ::= null
```

Zero valued; the **arglist** is empty.

2.5.22 Parabolic drive

```
<drive_caller> ::= parabolic ,
  <const_coef> ,
  <linear_coef> ,
  <parabolic_coef>
```

The function

$$\begin{aligned}
 f(t) = & \text{const_coef} \\
 & + \text{linear_coef} \cdot t \\
 & + \text{parabolic_coef} \cdot t^2
 \end{aligned}
 \tag{2.20}$$

2.5.23 Periodic drive

```
<drive_caller> ::= periodic ,
  <initial_time> ,
  <period> ,
  (DriveCaller)<func_drive>
```

The function

$$\begin{aligned}
 f(t) = 0 & & t < \text{initial_time} \\
 f(t) = \text{func_drive} \left(t - \text{initial_time} - \text{period} \cdot \text{floor} \left(\frac{t - \text{initial_time}}{\text{period}} \right) \right) & & t \geq \text{initial_time}
 \end{aligned}
 \tag{2.21}$$

2.5.24 Piecewise linear drive

```
<drive_caller> ::= piecewise linear ,
  <num_points> ,
  <point> , <value>
  [ , ... ]
```

Piecewise linear function; the first and the last point/value pairs are extrapolated in case a value beyond the extremes is required. Linear interpolation between pairs is used.

2.5.25 Ramp drive

```
<drive_caller> ::= ramp ,  
    <slope> ,  
    <initial_time> ,  
    { forever | <final_time> } ,  
    <initial_value>
```

The function

$$f(t) = \begin{cases} \text{initial_value} & t < \text{initial_time} \\ \text{slope} \cdot (t - \text{initial_time}) & \text{initial_time} \leq t \leq \text{final_time} \\ \text{slope} \cdot (\text{final_time} - \text{initial_time}) & \text{final_time} < t \end{cases} \quad (2.22)$$

2.5.26 Random drive

```
<drive_caller> ::= random ,  
    <amplitude_value> ,  
    <mean_value> ,  
    <initial_time> ,  
    { forever | <final_time> }  
    [ , steps , <steps_to_hold_value> ]  
    [ , seed , { time | <seed_value> } ]
```

the first optional entry, `steps_to_hold_value`, preceded by the keyword `steps`, sets the number of steps a random value must be held before generating a new random number. The second optional entry, preceded by the keyword `seed`, sets the new seed for the random number generator. A numeric `seed_value` value can be used, otherwise the keyword `time` uses the current time from the internal clock. A given seed can be used to ensure that two simulations use exactly the same random sequence (concurrent settings are not managed, so it is not very reliable).

2.5.27 Sample and hold drive

```
<drive_caller> ::= sample and hold ,  
    (DriveCaller) <function> ,  
    (DriveCaller) <trigger>  
    [ , initial value , <initial_value> ]
```

When `trigger` is non-zero, the value of `function` is recorded and returned; when `trigger` is zero, the last recorded value is returned. When the optional keyword `initial value` is present, if the trigger is initially zero, the value `initial_value` is returned until `trigger` becomes non-zero.

2.5.28 Scalar function drive

```
<drive_caller> ::= scalar function , " <scalar_function_name> "  
    [ , <scalar_function_definition> ]
```

The scalar function called `scalar_function_name` must exist; alternatively, it is defined inline according to `scalar_function_definition`.

Example.

```
scalar function: "this_is_a_test",
  multilinear,
    -20., 1.,
    -10., 0.1,
    0., 0.,
    10., 0.1,
    20., 1.;
set: integer MY_DRIVE = 1000;
drive caller: MY_DRIVE, scalar function, "this_is_a_test";
```

Alternative syntax example:

```
set: integer MY_DRIVE = 1000;
drive caller: MY_DRIVE, scalar function, "this_is_a_test",
  multilinear,
    -20., 1.,
    -10., 0.1,
    0., 0.,
    10., 0.1,
    20., 1.;
```

2.5.29 Sine drive

```
<drive_caller> ::= sine ,
  <initial_time> ,
  <angular_velocity> ,
  <amplitude> ,
  { [ - ] <number_of_cycles> | half | one | forever } ,
  <initial_value>
```

where `angular_velocity` is $2\pi/T$. The value of `number_of_cycles` determines the behavior of the drive. If it is positive, `number_of_cycles` - 1/2 oscillations are performed. If it is negative, the oscillations end after `number_of_cycles` - 3/4 cycles at the top of the sine, with null tangent. Special keywords can be used for `number_of_cycles`:

- **forever**: the oscillation never stops;
- **one**: exactly half period is performed;
- **half**: exactly a quarter of period is performed, so the function stops at

$$f(t) = \text{initial_value} \quad (2.23)$$

$$+ \begin{cases} 0 & t < \text{initial_time} \\ \text{amplitude} \cdot \sin(\text{angular_velocity} \cdot (t - \text{initial_time})) & \text{initial_time} \leq t \leq \text{initial_time} + \frac{\pi}{2 \cdot \text{angular_velocity}} \\ \text{amplitude} & \text{initial_time} + \frac{\pi}{2 \cdot \text{angular_velocity}} \leq t \end{cases} .$$

2.5.30 Step drive

```
<drive_caller> ::= step ,  
    <initial_time> ,  
    <step_value> ,  
    <initial_value>
```

The function

$$f(t) = \text{initial_value} + \begin{cases} 0 & t < \text{initial_time} \\ \text{step_value} & \text{initial_time} \leq t \end{cases} \quad (2.24)$$

2.5.31 String drive

```
<drive_caller> ::= string ,  
    " <expression_string> "
```

`expression_string` is a string, delimited by double quotes. It is parsed by the math parser every time the `drive` is invoked. The variable `Time` is kept up to date and can be used in the string to compute the return value. Another variable, `Var`, is set to the value provided by the caller in case the drive is called with an explicit argument as, for instance, in the `dof drive` (see Section 2.5.7); e.g.:

```
..., string, "e^(-Time)*cos(2.*pi*Time)" ...
```

generates a cosine modulated by an exponential.

Note: currently, the expression is parsed each time it is evaluated; this is not very efficient. A future version should build a tree-like form of the parsed expression, where non-symbolic branches are evaluated once for all, to speed up evaluation.

2.5.32 Tanh drive

```
<drive_caller> ::= tanh ,  
    <initial_time> ,  
    <amplitude> ,  
    <slope> ,  
    <initial_value>
```

This drive computes a function of the type

$$f(t) = \text{initial_value} + \text{amplitude} \cdot \tanh(\text{slope} \cdot (t - \text{initial_time})).$$

It is differentiable.

2.5.33 Time drive

```
<drive_caller> ::= time
```

Yields the current time; the arglist is empty.

2.5.34 Timestep drive

```
<drive_caller> ::= timestep
```

Yields the current timestep; the arglist is empty.

2.5.35 Unit drive

```
<drive_caller> ::= unit
```

Always 1; the `arglist` is empty.

2.5.36 Deprecated drive callers

- `one`; use `unit` instead.

2.5.37 Hints

Hints have been sponsored by Hutchinson CdR.

In some cases, during the analysis, different entities can be reinitialized as a consequence of some event, typically triggered by the waking up of a so-called `driven` element (See Section 8.17.2). The `DriveCaller` is re-parsed when the entity that owns it is sent a `hint` of the form

```
"drive{ (DriveCaller) <drive_caller> }"
```

is used, where `drive_caller` is an arbitrary `DriveCaller` specification. The whole hint argument needs to be enclosed in double quotes because MBDyn treats it as a string. It is not parsed until needed. At that point, the string (after removing the enclosing double quotes) is fed into the parser.

Typically, the use of a `hint` is needed when the specification of the drive caller parameters depends on the configuration of the system when the entity that uses it is waken up. For example, a `distance joint`, an element that enforces the distance between two nodes, may be waken up by some event. In the example reported below, the initial value of the `cosine` drive caller is computed by extracting the current distance between the nodes at the time the element is waken up:

```
driven: 1, string, "Time > 10.",
      hint, "drive{cosine, 10., 2., .25, half, model::distance_prev(10, 20)}",
joint: 1, distance,
      10,
      20,
      const, 1; # this length is not actually used,
                # since the element is inactive until "Time > 10.";
                # subsequently, the length is replaced by the hint
```

Note that the function `distance_prev` (from the built-in `model` namespace) is used instead of the function `distance`. This is needed because hints are parsed *after* the prediction of the state of the model for a subsequent new time step, i.e. they extract information from the model *after* its current state has been modified by the prediction for the next time step. As a consequence, unless an exactly steady state was reached for a sufficient number of steps to result in a prediction that does not change the state of that part of model, unexpected valued could result from a call to a function that returns information based on the *current* state of the model. On the contrary, to access the last converged state of the model within a `hint` one needs to call `mode` namespace functions that refer to the *previous* time step. See Section 2.1.4 for further details on the functions of the `model` namespace.

2.5.38 Template Drive

A particular `DriveCaller` is the template drive caller, indicated as `TplDriveCaller<Entity>`. It basically consists in a drive with dimensionality other than scalar. The family of the `TplDriveCaller<Entity>` object is very large. The type of the `TplDriveCaller<Entity>` is declared as

```
<tpl_drive_caller> ::=
  { <tpl_drive_caller_type> [ , <arglist> ]
    | reference , <label> }
```

where `arglist`, if any, is a comma-separated list of arguments that depends on `tpl_drive_caller_type`.

If the alternative format is used, the keyword `reference` must be followed by the label of an already defined, valid drive caller (See Section 2.3.12).

There are four types of template drive callers:

- the `null` template drive caller:

```
<tpl_drive_caller> ::= null
```

a shortcut for nothing;

- the `single` template drive caller:

```
<tpl_drive_caller> ::= single ,
  <entity> , (DriveCaller) <drive_caller>
```

corresponding to `entity × drive_caller`, where `entity` is a constant of the expected type (scalar, 3×1 vector, 6×1 vector are the types currently defined, but, since a C++ template has been used, the implementation of new ones is straightforward);

- the `component` template drive caller:

```
<tpl_drive_caller> ::= component , [ <type> , ]
  <component_drive_caller>
  [ , ... ]
  # as many component_drive_caller as the dimensionality of entity

<type> ::= { sym | diag }
```

where each element consists of the corresponding drive caller; the optional `type` is used with matrix-type entities to specify special matrix layouts:

- `sym` indicates that only the upper triangular components are expected;
- `diag` indicates that only the diagonal components are expected;

- the `array` template drive caller:

```
<component_drive_caller> ::= { inactive | (DriveCaller) <drive_caller> }

<tpl_drive_caller> ::= array ,
  <num_template_drive_callers> ,
  <tpl_drive_caller>
  [ , ... ]
  # num_template_drive_callers instances of tpl_drive_caller
```

which linearly combines a set of arbitrary template drive callers.

In case of scalar values, a normal drive caller is actually instantiated, so no overhead is added.

In the case of the **component** template drive caller, as many instances of a **DriveCaller** as the dimensionality of the template drive caller are expected.

At least one **tpl_drive_caller** is expected in the case of the **array** template drive caller. If **num_template_drive_callers** is exactly 1, only a single template drive caller is actually constructed, thus avoiding the overhead related to handling the template drive caller array.

The template drive can be re-parsed as a consequence of a **hint** (see Section 2.5.37). In this case, the syntax of the **hint** is:

```
"drive<n>{ <tpl_drive_caller> }"
```

where **n** is the dimension of **entity**. For example, a **TplDriveCaller<Vec3>** drive hint is:

```
"drive3{ 1., 0., 0., sine, 0., 2*pi, 1., forever, 0.}"
```

which corresponds to a vector 1., 0., 0. that multiplies a **sine** drive caller with the parameters illustrated above.

2.6 Scalar functions

A **ScalarFunction** object allows to compute the value of a function. Almost every scalar function is of type **DifferentiableScalarFunction**, derived from **ScalarFunction**, and allows to compute the derivatives of the function as well. Currently implemented scalar functions are

- **const**: $f(x) = c$
- **exp**: $f(x) = m \cdot b^{c \cdot x}$
- **log**: $f(x) = m \cdot \log_b(c \cdot x)$
- **pow**: $f(x) = x^p$
- **linear**: linear interpolation between the two points (x_1, y_1) and (x_2, y_2)
- **cubicspline**: cubic natural spline interpolation between the set of points $\{(x_i, y_i), i \in [1, k \geq 3]\}$
- **multilinear**: multilinear interpolation between the set of points $\{(x_i, y_i), i \in [1, k \geq 2]\}$
- **chebychev**: Chebychev interpolation between the set of points $\{a, b\}$
- **sum**: $f(x) = f_1(x) + f_2(x)$
- **sub**: $f(x) = f_1(x) - f_2(x)$
- **mul**: $f(x) = f_1(x) \cdot f_2(x)$
- **div**: $f(x) = f_1(x)/f_2(x)$

Every **ScalarFunction** card follows the format

```
<card> ::= scalar function : " <unique_scalar_func_name> " ,  
      <scalar_func_type> ,  
      <scalar_func_args>
```

The name of the scalar function is a string, thus it needs to be enclosed in double quotes.

The type of scalar function, **scalar_func_type**, together with relevant arguments, **scalar_func_args**, are as follows:

Const Scalar Function

```
<scalar_func_type> ::= const  
  
<scalar_func_args> ::= <const_coef>
```

Note: if the `scalar_func_type` is omitted, a `const` scalar function is assumed.

Example.

```
scalar function: "const_function", const, 1.e-2;
```

Exp Scalar Function

```
<scalar_func_type> ::= exp  
  
<scalar_func_args> ::=  
  [ base , <base> , ]  
  [ coefficient , <coef> , ]  
  <multiplier_coef>
```

$$f(x) = \text{multiplier_coef} \cdot \text{base}^{\text{coef} \cdot x} \quad (2.25)$$

Note: the optional `base` must be positive (defaults to 'e', Neper's number); `coef` defaults to 1.

Example.

```
scalar function: "exp_function", exp, 1.e-2; # 1.e-2*e^x
```

Log Scalar Function

```
<scalar_func_type> ::= log  
  
<scalar_func_args> ::=  
  [ base , <base> , ]  
  [ coefficient , <coef> , ]  
  <multiplier_coef>
```

$$f(x) = \text{multiplier_coef} \cdot \log_{\text{base}}(\text{coef} \cdot x) \quad (2.26)$$

Note: the optional `base` (defaults to 'e', Neper's number, resulting in natural logarithms) must be positive. The optional value `coef`, prepended by the keyword `coefficient`, (defaults to 1). The argument must be positive, otherwise an exception is thrown.

Example.

```
scalar function: "log_function", log, 1.e-2; # 1.e-2*log(x)
```

Pow Scalar Function

```
<scalar_func_type> ::= pow
<scalar_func_args> ::= <exponent_coef>
```

$$f(x) = x^{\text{exponent_coef}} \quad (2.27)$$

Example.

```
scalar function: "pow_function", pow, 3.; # x^3.
```

Linear Scalar Function

```
<scalar_func_type> ::= linear
<scalar_func_args> ::= <point> , <point>
<point>                ::= <x> , <y>
```

A line passing through the two points,

$$f(x) = y_i + \frac{y_f - y_i}{x_f - x_i} (x - x_i) \quad (2.28)$$

where x_i, y_i are the coordinates of the first point, while x_f, y_f are the coordinates of the second point.

Example.

```
scalar function: "linear_function", linear, 0., 0., 1., 1.;
```

Cubic Natural Spline Scalar Function

```
<scalar_func_type> ::= cubic spline
<scalar_func_args> ::= [ do not extrapolate , ]
    <point> ,
    <point>
    [ , ... ]
    [ , end ]
<point>                ::= <x> , <y>
```

The `end` delimiter is required if the card needs to continue (i.e. the `ScalarFunction` definition is embedded in a more complex statement); it can be omitted if the card ends with a semicolon right after the last point.

Multilinear Scalar Function

```
<scalar_func_type> ::= multilinear

<scalar_func_args> ::= [ do not extrapolate , ]
    <point> ,
    <point>
    [ , ... ]
    [ , end ]

<point>                ::= <x> , <y>
```

Unless `do not extrapolate` is set, when the input is outside the provided values of `x` the value is extrapolated using the slope of the nearest point pair.

Chebyshev Scalar Function

```
<scalar_func_type> ::= chebychev

<scalar_func_args> ::=
    <lower_bound> , <upper_bound> ,
    [ do not extrapolate , ]
    <coef_0> [ , <coef_1> [ , ... ] ]
    [ , end ]
```

Chebyshev polynomials of the first kind are defined as

$$T_n(\xi) = \cos(n \cos^{-1}(\xi)), \quad (2.29)$$

with

$$\xi = 2 \frac{x}{b-a} - \frac{b+a}{b-a}, \quad (2.30)$$

where $a = \text{lower_bound}$ and $b = \text{upper_bound}$, which corresponds to the series

$$T_0(\xi) = 1 \quad (2.31)$$

$$T_1(\xi) = \xi \quad (2.32)$$

...

$$T_n(\xi) = 2\xi T_{n-1}(\xi) - T_{n-2}(\xi). \quad (2.33)$$

For example, the first five coefficients are

$$T_0(\xi) = 1 \quad (2.34)$$

$$T_1(\xi) = \xi \quad (2.35)$$

$$T_2(\xi) = 2\xi^2 - 1 \quad (2.36)$$

$$T_3(\xi) = 4\xi^3 - 3\xi \quad (2.37)$$

$$T_4(\xi) = 8\xi^4 - 8\xi^2 + 1 \quad (2.38)$$

This scalar function implements the truncated series in the form

$$f(x) = \sum_{k=0,n} c_k T_k(\xi), \quad (2.39)$$

where $c_k = \text{coef_} \langle k \rangle$. The first derivative of the series is obtained by considering

$$\frac{d}{d\xi} T_0(\xi) = 0 \tag{2.40}$$

$$\frac{d}{d\xi} T_1(\xi) = 1 \tag{2.41}$$

...

$$\frac{d}{d\xi} T_n(\xi) = 2T_{n-1}(\xi) + 2\xi \frac{d}{d\xi} T_{n-1}(\xi) - \frac{d}{d\xi} T_{n-2}(\xi), \tag{2.42}$$

so the first derivative of the scalar function is

$$\frac{d}{dx} f(x) = \frac{d\xi}{dx} \sum_{k=1,n} c_k \frac{d}{d\xi} T_k(\xi). \tag{2.43}$$

Subsequent derivatives follow the rule

$$\frac{d^i}{d\xi^i} T_n(\xi) = 2i \frac{d^{i-1}}{d\xi^{i-1}} T_{n-1}(\xi) + 2\xi \frac{d^i}{d\xi^i} T_{n-1}(\xi) - \frac{d^i}{d\xi^i} T_{n-2}(\xi). \tag{2.44}$$

Differentiation of order higher than 1 is not currently implemented.

Sum Scalar Function

```
<scalar_func_type> ::= sum
```

```
<scalar_func_args> ::= (ScalarFunction) <f1> , (ScalarFunction) <f2>
```

$$f(x) = f1(x) + f2(x) \tag{2.45}$$

Example.

```
scalar function: "first_function", const, 1.;
scalar function: "second_function", const, 2.;
scalar function: "sum_function", sum, "first_function", "second_function";
```

Sub Scalar Function

```
<scalar_func_type> ::= sub
```

```
<scalar_func_args> ::= (ScalarFunction) <f1> , (ScalarFunction) <f2>
```

$$f(x) = f1(x) - f2(x) \tag{2.46}$$

Example.

```
scalar function: "first_function", const, 1.;
scalar function: "second_function", const, 2.;
scalar function: "sub_function", sub, "first_function", "second_function";
```

Mul Scalar Function

```
<scalar_func_type> ::= mul
```

```
<scalar_func_args> ::= (ScalarFunction) <f1> , (ScalarFunction) <f2>
```

$$f(x) = f1(x) \cdot f2(x) \tag{2.47}$$

Example.

```
scalar function: "first_function", const, 1.;  
scalar function: "second_function", const, 2.;  
scalar function: "mul_function", mul, "first_function", "second_function";
```

Div Scalar Function

```
<scalar_func_type> ::= div
```

```
<scalar_func_args> ::= (ScalarFunction) <f1> , (ScalarFunction) <f2>
```

$$f(x) = f1(x) / f2(x) \tag{2.48}$$

Note: division by zero is checked, and an exception is thrown.

Example.

```
scalar function: "first_function", const, 1.;  
scalar function: "second_function", const, 2.;  
scalar function: "div_function", div, "first_function", "second_function";
```

2.7 Friction

A friction-based element needs the definition of at least a friction model and of a shape function.

2.7.1 Friction models

The `friction_model` input format is:

```
<friction_model> ::= <friction_type> ,  
                  <friction_arglist>
```

where `friction_function` gives the static friction as a function of sliding velocity.

Currently implemented friction models are:

1. `modlugre`

This friction model is based on the one presented in [2] by Pierre Dupont, Vincent Hayward, Brian Armstrong and Friedhelm Altpeter, known as ‘LuGre’. The input format is:

```

<friction_model> ::= modlugre

<friction_arglist> ::=
  <sigma0> ,
  <sigma1> ,
  <sigma2> ,
  <kappa> ,
  (ScalarFunction) <friction_function>

```

2. discrete coulomb

This is a Coulomb model with viscous friction and internal states to resolve stick/slip conditions.

```

<friction_model> ::= discrete coulomb

<friction_arglist> ::=
  [ sigma2 , <sigma2> , ]
  [ velocity ratio , <vel_ratio> , ]
  (ScalarFunction) <friction_function>

```

where `sigma2` gives the viscous friction; `vel_ratio` defaults to 0.8, and is used to discriminate stick/slip conditions. Note that `discrete coulomb` cannot handle stick conditions when the reaction force in the constraint goes to zero (resulting in a singular Jacobian matrix). The `preload` parameter needs to be used to make sure the joint is preloaded as appropriate (see the documentation of the specific joint for details).

2.7.2 Shape functions

The input format of shape functions is:

```

<shape_function> ::= <shape_function_type> ,
  <shape_function_arglist>

```

Currently implemented shape functions for friction models are:

1. simple

This shape function is equal to one. It does not need arguments.

```

<shape_function_type> ::= simple

<shape_function_arglist> ::=

```

2. simple plane hinge

```

<shape_function_type> ::= simple plane hinge

<shape_function_arglist> ::= <radius>

```

This is the shape function of a `revolute hinge` with radius equal to `radius` and subject to small loads.

2.8 Shapes

The **Shape** entities are objects that return a value depending on one (or two, for 2D shapes) dimensionless abscissa, ranging $[-1, 1]$. At present, only 1D shapes are used, by aerodynamic elements. A **Shape** input format is:

```
<shape_1D> ::= <shape_type> ,
           <shape_arglist>
```

The shapes currently available are:

1. **const**

```
<shape_type> ::= const
<shape_arglist> ::= <const_value>
```

$$f = \text{const_value} \quad (2.49)$$

2. **piecewise const**

```
<shape_type> ::= piecewise const
<shape_arglist> ::=
  <number_of_points> ,
  <abscissa> , <value>
  [ , ... ]
```

$$f(\xi) = \text{value}_1 \quad \xi \leq \text{abscissa}_1 \quad (2.50a)$$

$$f(\xi) = \text{value}_i \quad \text{abscissa}_i \leq \xi < \text{abscissa}_{i+1} \quad (2.50b)$$

$$f(\xi) = \text{value}_N \quad \text{abscissa}_N \leq \xi \quad (2.50c)$$

3. **linear**

```
<shape_type> ::= linear
<shape_arglist> ::=
  <value_at_-1> ,
  <value_at_1>
```

$$f(\xi) = \frac{\text{value_at_} -1 + \text{value_at_} 1}{2} + \frac{\text{value_at_} -1 - \text{value_at_} 1}{2} \xi \quad (2.51)$$

4. **piecewise linear**

```
<shape_type> ::= piecewise linear
<shape_arglist> ::=
  <number_of_points> ,
  <abscissa> , <value>
  [ , ... ]
```

$$f(\xi) = \text{value}_1 \quad \xi \leq \text{abscissa}_1 \quad (2.52a)$$

$$f(\xi) = \text{value}_i \frac{\text{abscissa}_{i+1} - \xi}{\text{abscissa}_{i+1} - \text{abscissa}_i} + \text{value}_{i+1} \frac{\xi - \text{abscissa}_i}{\text{abscissa}_{i+1} - \text{abscissa}_i} \quad \text{abscissa}_i \leq \xi < \text{abscissa}_{i+1} \quad (2.52b)$$

$$f(\xi) = \text{value}_N \quad \text{abscissa}_N \leq \xi \quad (2.52c)$$

5. parabolic

```
<shape_type> ::= parabolic
```

```
<shape_arglist> ::=
    <value_at_-1> ,
    <value_at_0> ,
    <value_at_1>
```

$$f(\xi) = \text{value_at_} -1 \cdot \xi (\xi - 1) + \text{value_at_} 0 \cdot (1 - \xi^2) + \text{value_at_} 1 \cdot \xi (\xi + 1) \quad (2.53)$$

This form of input has been chosen since, being the shapes mainly used to interpolate values, it looks more “natural” to insert the mapping values at characteristic points. For **piecewise linear** shapes, there must be **number_of_points** pairs of abscissæ and values; abscissæ must be in the range $[-1, 1]$, in strict ascending order.

2.9 Constitutive Laws

Every time a “deformable” entity requires a constitutive law, a template constitutive law is read. This has been implemented by means of C++ templates in order to allow the definition of a general constitutive law when possible. The “deformable” elements at present are:

- **rod** and **genel spring** and related elements (1D);
- **deformable hinge** and **deformable displacement joint** elements (3D);
- **deformable joint** and **beam** elements (6D).

Constitutive laws are also used in non-structural components, to allow some degree of generality in defining input/output relationships. Some constitutive laws are meaningful only when related to some precise dimensionality. In some special cases, general purpose (**genel**) elements use 1D constitutive laws to express an arbitrary dependence of some value on a scalar state of the system. Table 2.7 shows the availability of each constitutive law.

The meaning of the input and output parameters of a constitutive law is dictated by the entity that uses it. In general, the user should refer to the element the constitutive law is being instantiated for in order to understand what the input and the output parameters are supposed to be.

Usually, constitutive laws can be directly defined when required, according to the definition of an element. However, the special card described in Section 2.3.1 allows to define constitutive laws stand-alone, and attach them to the elements by means of the following mechanism:

```
<constitutive_law> ::=
    { <constitutive_law_definition> | reference , <label> }
```

Table 2.7: Constitutive laws dimensionality

Constitutive law	1D	3D	6D
linear elastic, linear elastic isotropic	✓	✓	✓
linear elastic generic	✓	✓	✓
linear elastic generic axial torsion coupling			✓
cubic elastic generic	✓	✓	
log elastic	✓		
linear elastic bistop	✓	✓	✓
double linear elastic	✓	✓	
isotropic hardening elastic	✓	✓	✓
scalar function elastic, scalar function elastic isotropic	✓	✓	✓
scalar function elastic orthotropic	✓	✓	✓
linear viscous, linear viscous isotropic	✓	✓	✓
linear viscous generic	✓	✓	✓
linear viscoelastic, linear viscoelastic isotropic	✓	✓	✓
linear viscoelastic generic	✓	✓	✓
linear time variant viscoelastic generic	✓	✓	✓
linear viscoelastic generic axial torsion coupling			✓
cubic viscoelastic generic	✓	✓	
double linear viscoelastic	✓	✓	
turbulent viscoelastic	✓		
linear viscoelastic bistop	✓	✓	✓
shock absorber	✓		
symbolic elastic	✓	✓	✓
symbolic viscous	✓	✓	✓
symbolic viscoelastic	✓	✓	✓
ann elastic	✓	✓	✓
ann viscoelastic	✓	✓	✓
nlsf elastic	✓	✓	✓
nlsf viscous	✓	✓	✓
nlsf viscoelastic	✓	✓	✓
nlp elastic	✓	✓	✓
nlp viscous	✓	✓	✓
nlp viscoelastic	✓	✓	✓
Constitutive law wrapper	1D	3D	6D
array (wrapper)	✓	✓	✓
bistop (wrapper)	✓	✓	✓
invariant angular (wrapper)		✓	

where `<constitutive_law_definition>` is described in the following, while `<label>` is the label of a previously defined constitutive law of the appropriate dimensionality, as described in Section 2.3.1.

The constitutive laws are entered as follows:

```
<constitutive_law_definition> ::= <specific_const_law>
    [ , prestress , (Entity) <prestress> ]
    [ , prestrain , (TplDriveCaller<Entity>) <prestrain> ]

<specific_const_law> ::= <const_law_name> ,
    <const_law_data>
```

where `const_law_name` is the name of the constitutive law and `const_law_data` depends on the specific constitutive law. The latter fields, whose type depends on the dimensionality of the constitutive law, are optional, under the assumption that the constitutive law is the last portion of a card, or that any ambiguity can be avoided. The data specific to the currently available constitutive laws must be entered as follows:

2.9.1 Linear elastic, linear elastic isotropic

```
<specific_const_law> ::= linear elastic [ isotropic ] ,
    (Scalar) <stiffness>
```

the isotropic stiffness coefficient; the word `isotropic` can be omitted, essentially because it has no meaning for scalar constitutive laws.

Example.

```
constitutive law: 1, name, "scalar isotropic law",
    1, linear elastic, 1.e9;
constitutive law: 2, name, "3D isotropic law",
    3, linear elastic isotropic, 1.e9;
constitutive law: 3, name, "6D isotropic law",
    6, linear elastic isotropic, 1.e9;
```

2.9.2 Linear elastic generic

```
<specific_const_law> ::= linear elastic generic ,
    (DerivativeOfEntity) <stiffness>
```

the stiffness matrix. In case of 1D, the type is scalar, and there is no distinction between `generic` and `isotropic`, while, in case of $N \times 1$ vectors, the type is the corresponding $N \times N$ matrix.

Example.

```
constitutive law: 1, name, "scalar isotropic law",
    1, linear elastic generic, 1.e9;
constitutive law: 2, name, "3D isotropic law",
    3, linear elastic generic,
    sym, 1.e9, 0., 0.,
    1.e6, -1.e5,
    1.e6;
```

```

constitutive law: 3, name, "6D isotropic law",
  6, linear elastic generic,
    diag, 1.e9, 1.e9, 1.e9, 1.e6, 1.e6, 1.e6;

```

2.9.3 Linear elastic generic axial torsion coupling

```

<specific_const_law> ::=
  linear elastic generic axial torsion coupling ,
    (DerivativeOfEntity) <stiffness> ,
    (Scalar) <coupling_coefficient>

```

this is defined only for 6×1 vectors, where the torsion stiffness, coefficient a_{44} in the stiffness matrix, depends linearly on the axial strain, ε_1 , by means of `coupling_coefficient`, i.e. the run-time torsion stiffness is

$$a_{44} = GJ + \text{coupling_coefficient} \cdot \varepsilon_1. \quad (2.54)$$

This `<coupling_coefficient>`, in the classical nonlinear beam theory, is estimated according to some geometric property [3]; a general approach to the computation of prestressed beam properties is presented in [4], which is implemented in some versions of the ANBA software.

2.9.4 Cubic elastic generic

```

<specific_const_law> ::=
  cubic elastic generic
    (Entity) <stiffness_1> ,
    (Entity) <stiffness_2> ,
    (Entity) <stiffness_3>

```

this is defined only for scalar and 3×1 vectors; the constitutive law is written according to the formula

$$f = \text{stiffness}_1 \cdot \varepsilon + \text{stiffness}_2 \cdot |\varepsilon| \varepsilon + \text{stiffness}_3 \cdot \varepsilon^3 \quad (2.55)$$

2.9.5 Log elastic

```

<specific_const_law> ::= log elastic ,
  (DerivativeOfEntity) <stiffness>

```

this is defined only for scalars. The force is defined as:

$$f = \text{stiffness} \cdot \log(1 + \varepsilon) \quad (2.56)$$

2.9.6 Linear elastic bistop

```

<specific_const_law> ::= linear elastic bistop ,
  (DerivativeOfEntity) <stiffness> ,
  [ initial status , { inactive | active | (bool) <status> } , ]
  (DriveCaller) <activating_condition> ,
  (DriveCaller) <deactivating_condition>

```

2.9.7 Double linear elastic

```
<specific_const_law> ::= double linear elastic ,  
  (Scalar) <stiffness_1> ,  
  (Scalar) <upper_strain> ,  
  (Scalar) <lower_strain> ,  
  (Scalar) <stiffness_2>
```

this is defined for scalar and 3×1 vectors. In the scalar case the meaning of the entries is straightforward, while in case of 3×1 vectors the constitutive law is isotropic but in the local direction 3, where, in case of strain out of the upper or lower bound, the `stiffness_2` is used.

2.9.8 Isotropic hardening elastic

```
<specific_const_law> ::= isotropic hardening elastic ,  
  (Scalar) <stiffness> ,  
  (Scalar) <reference_strain>  
  [ , linear stiffness , <linear_stiffness> ]
```

this constitutive law is defined as follows:

$$f = \text{stiffness} \frac{\beta + \alpha |\varepsilon|^2}{1 + \alpha |\varepsilon|^2} \varepsilon$$

where $\alpha = 3/|\text{reference_strain}|^2$, and $\beta = \text{linear_stiffness}/\text{stiffness}$. The resulting constitutive law, in the scalar case, is somewhat soft/hard when β is lower/greater than 1 and $|\varepsilon|$ is smaller than `reference_strain`, approaching `linear_stiffness` when $\varepsilon \rightarrow 0$, while it grows to quasi-linear for larger $|\varepsilon|$, with slope `stiffness`.

2.9.9 Scalar function elastic, scalar function elastic isotropic

This constitutive law is based on a `ScalarFunction` (Section 2.6) to represent an analytical force-displacement curve of a single variable that is automatically differentiated to compute the slope of the curve, namely the local stiffness.

```
<specific_const_law> ::= scalar function elastic [ isotropic ] ,  
  (DifferentiableScalarFunction) <function>
```

this constitutive law is defined as follows:

$$f_i = \text{function}(\varepsilon_i)$$

the force is computed for each direction as a function of the respective strain component using the same function. When used for 1D elements, the word `isotropic` can be omitted.

2.9.10 Scalar function elastic orthotropic

```
<specific_const_law> ::= scalar function elastic [ orthotropic ] ,  
  { (DifferentiableScalarFunction) <function> | null }  
  [ , ... ]
```

this constitutive law is defined as follows:

$$f_i = \text{function}_i(\varepsilon_i)$$

the force is computed for each direction as a function of the respective strain component using a specific function for each component. If no force ought to be used for a direction, for example because that direction is constrained by a kinematic joint, the keyword `null` can be used to indicate that no function is expected for that component. When used for 1D elements, the word `orthotropic` can be omitted; note that in this case a `scalar function elastic isotropic` constitutive law is actually instantiated.

Example.

```
scalar function: "myfunc", multilinear,
  -1., -100.,
  -.5, -70.,
  0., 0.,
  .5, 70.,
  1., 100.;
constitutive law: 1000, 3,
  scalar function elastic orthotropic,
  null,
  null,
  "myfunc";
```

indicates that the constitutive law is only defined in direction 3 as a multilinear function.

2.9.11 Linear viscous, linear viscous isotropic

```
<specific_const_law> ::= linear viscous [ isotropic ] ,
  (Scalar) <viscosity>
```

the linear viscous coefficient.

Note: this constitutive law does not require any prestrain template drive caller.

2.9.12 Linear viscous generic

```
<specific_const_law> ::= linear viscous generic ,
  (DerivativeOfEntity) <viscosity>
```

the linear viscous $N \times N$ matrix.

Note: this constitutive law does not require any prestrain template drive caller.

2.9.13 Linear viscoelastic, linear viscoelastic isotropic

```
<specific_const_law> ::= linear viscoelastic [ isotropic ] ,
  (Scalar) <stiffness> ,
  { (Scalar) <viscosity>
    | proportional , (Scalar) <factor> }
```

the isotropic stiffness and viscosity coefficients. If `proportional` is given, then `viscosity = factor · stiffness`.

2.9.14 Linear viscoelastic generic

```
<specific_const_law> ::= linear viscoelastic generic ,  
  (DerivativeOfEntity) stiffness ,  
  { (DerivativeOfEntity) <viscosity>  
    | proportional , (Scalar) <factor> }
```

the linear stiffness and viscosity $N \times N$ matrices. If `proportional` is given, then `viscosity = factor · stiffness`.

2.9.15 Linear time variant viscoelastic generic

```
<specific_const_law> ::= linear time variant viscoelastic generic ,  
  (DerivativeOfEntity) <stiffness> ,  
  (DriveCaller) <stiffness_scale> ,  
  { (DerivativeOfEntity) <viscosity>  
    | proportional , (Scalar) <factor> } ,  
  (DriveCaller) <viscosity_scale>
```

the linear stiffness and viscosity matrices are multiplied by the respective scale factors,

$$\mathbf{f} = \text{stiffness} \cdot \text{stiffness_scale} \cdot \boldsymbol{\varepsilon} + \text{viscosity} \cdot \text{viscosity_scale} \cdot \dot{\boldsymbol{\varepsilon}}. \quad (2.57)$$

If `proportional` is given, then `viscosity = factor · stiffness`.

Example.

```
linear time variant visco elastic,  
  1000., cosine, 2., pi/.2, .1/2, half, 1.,  
  100., cosine, 2., pi/.2, 1/2, half, 1.
```

At 2s, the stiffness grows of 10% from the nominal value in .2s, while the damping doubles. See Section 2.5 for details on the syntax of `DriveCaller` entities.

Beware that arbitrarily changing the stiffness and the damping of an elastic component during the execution of the simulation may have no physical meaning. The intended use of this feature is for tailoring the analysis; for example, a higher damping level may be desirable to smooth out a non-physical transient, and later return to the appropriate damping value.

2.9.16 Linear viscoelastic generic axial torsion coupling

```
<specific_const_law> ::=  
  linear viscoelastic generic axial torsion coupling ,  
  (DerivativeOfEntity) <stiffness> ,  
  { (DerivativeOfEntity) <viscosity>  
    | proportional , (Scalar) <factor> }  
  (Scalar) <coupling_coefficient>
```

this is defined only for 6×1 vectors; it is the viscoelastic extension of the `linear elastic generic axial torsion coupling` constitutive law (see Section 2.9.3).

2.9.17 Cubic viscoelastic generic

```
<specific_const_law> ::=  
  cubic elastic generic  
    (Entity) <stiffness_1> ,  
    (Entity) <stiffness_2> ,  
    (Entity) <stiffness_3> ,  
    (DerivativeOfEntity) <viscosity>
```

this is defined only for scalar and 3×1 vectors; the constitutive law is written according to the formula

$$f = \text{stiffness_1} \cdot \varepsilon + \text{stiffness_2} \cdot |\varepsilon| \varepsilon + \text{stiffness_3} \cdot \varepsilon^3 + \text{viscosity} \cdot \dot{\varepsilon}$$

and it is mainly intended for use with human body models where the stiffness of the joints is typically given in this form.

2.9.18 Double linear viscoelastic

```
<specific_const_law> ::= double linear viscoelastic ,  
  (Scalar) <stiffness_1> ,  
  (Scalar) <upper_strain> ,  
  (Scalar) <lower_strain> ,  
  (Scalar) <stiffness_2> ,  
  (Scalar) <viscosity>  
  [ , second damping , (Scalar) <viscosity_2> ]
```

this is analogous to the **double linear elastic** constitutive law, except for the isotropic viscosity term. The second viscosity value is used when the strain is outside the **lower_strain – upper_strain** range.

When this constitutive law is used with 3×1 vectors, the double linear elastic and viscous term only applies to component 3.

2.9.19 Turbulent viscoelastic

```
<specific_const_law> ::= turbulent viscoelastic ,  
  (Scalar) <stiffness> ,  
  (Scalar) <parabolic_viscosity>  
  [ , (Scalar) <threshold>  
    [ , (Scalar) <linear_viscosity> ] ]
```

the constitutive law has the form:

$$f = \text{stiffness} \varepsilon + k \dot{\varepsilon}$$

where:

$$k = \begin{cases} \text{linear_viscosity} & |\dot{\varepsilon}| \leq \text{threshold} \\ \text{parabolic_viscosity} & |\dot{\varepsilon}| > \text{threshold} \end{cases}$$

if **threshold** is null, or not defined, the constitutive law is always parabolic. If the **linear_viscosity** is not defined, it is computed based on **parabolic_viscosity** and on **threshold** to give a continuous force curve (with discontinuous slope). Otherwise, it can be set by the user to give a discontinuous force curve, as observed in some fluids at intermediate Reynolds number.

2.9.20 Linear viscoelastic bistop

```
<specific_const_law> ::= linear viscoelastic bistop ,
  (DerivativeOfEntity) <stiffness> ,
  (DerivativeOfEntity) <viscosity> ,
  [ initial status , { inactive | active | (bool) <status> } , ]
  (DriveCaller) <activating_condition> ,
  (DriveCaller) <deactivating_condition>
```

2.9.21 GRAALL damper

This is a very experimental constitutive law, based on a nonlinear model for a hydraulic damper to be used in landing gear modeling. Basically, it requires the user to supply the name of the GRAALL-style input file with damper data. It will be documented as soon as it reaches an appreciable level of stability. See also the `shock absorber` constitutive law.

2.9.22 shock absorber

This constitutive law implements a landing gear hydraulic shock absorber:

```
<specific_const_law> ::= shock absorber ,
  [ prestrain , <value> , ]
  <reference_pressure> ,
  <reference_area_for_force_computation> ,
  <interaction_coefficient> ,
  <polytropic_exponent> ,
  [ epsilon max , <upper_strain_bound> , ]
  [ epsilon min , <lower_strain_bound> , ]
  [ penalty , <penalty_factor_for_strain> ,
    <penalty_factor_for_strain_rate> , ]
  [ metering , <metering_area> ,
    [ negative , <metering_area_for_negative_strain_rate> , ]
  [ orifice , <orifice_area> , ]
  <fluid_area> ,
  <fluid_density> ,
  <drag_coefficient/_reference_length> # scales strain rate to velocity
  [ , friction , <reference_epsilon_prime> ,
    <friction_amplitude_coefficient> ]
```

where

- the `interaction_coefficient` is represented by

$$\text{kinematic scale} \cdot \frac{LA}{V_0}$$

where `kinematic scale` is the ratio between the stroke of the shock absorber and that of the gas;

- `upper_strain_bound` is the upper strain bound; it must be at least larger than the prestrain, and defaults to 0, i.e. the shock absorber, at rest, is assumed to be fully extended;
- `lower_strain_bound` is the lower strain bound; it must be at least smaller than the prestrain, and defaults to -0.5, i.e. the shock absorber is assumed to allow a contraction equal to half its full length;

- the `penalty_factor_for_strain` defaults to $1e+9$; it is active only when strain bounds are violated;
- the `penalty_factor_for_strain_rate` defaults to 0; it is active only when strain bounds are violated;
- the `metering_area` is given by a `DriveCaller` and is strain dependent; if the keyword `negative` is used, then the `metering_area_for_negative_strain_rate` is used when the strain rate is negative, i.e. the shock absorber is being compressed, while `metering_area` is used only when the shock absorber is extending;
- the `orifice` drive determines the area of an additional orifice, which essentially depends on the sign of the strain rate; it is used to implement relief valves;
- ...

This constitutive law contributes to the output of the element it is associated with. The name associated to each output contribution can be used to reference the corresponding value as element private data. The contributions are

1. "p" gas pressure;
2. "A" metering area;
3. "Fe" elastic force;
4. "Fv" viscous force.

2.9.23 symbolic elastic

The implementation of the family of `symbolic` constitutive laws is based on GiNaC (<http://www.ginac.de/>), a free software package for symbolic algebra manipulation. It is essentially used to automatically differentiate the user-supplied expression that describes the relationship between the output and the input.

The syntax is

```
<specific_const_law> ::= symbolic elastic ,
    epsilon , " <epsilon> " [ , ... ] ,
    expression , " <expression> " [ , ... ]
```

where `epsilon` is the symbol describing the input parameter as it will be used in `expression`.

For constitutive laws with more than one dimension, a string for each `epsilon` and one for each `expression` are expected. For example:

```
# 1D symbolic constitutive law
constitutive law: 1001, 1, symbolic elastic,
    epsilon, "eps",
    expression, "1000.*eps + 5.*eps^3";
# 3D symbolic constitutive law
constitutive law: 1003, 3, symbolic elastic,
    epsilon, "eps1", "eps2", "eps3",
    expression,
    "1000.*eps1 + 5.*eps1^3 - 10.*eps2*eps3",
    "1000.*eps2 + 5.*eps2^3 - 10.*eps3*eps1",
    "1000.*eps3 + 5.*eps3^3 - 10.*eps1*eps2";
```

Note: right now, the symbols defined within the mathematical parser are not available within symbolic constitutive laws.

2.9.24 symbolic viscous

The syntax is

```
<specific_const_law> ::= symbolic viscous ,  
    epsilon_prime , " <epsilon_prime> " [ , ... ] ,  
    expression , " <expression> " [ , ... ]
```

where `epsilon_prime` is the symbol describing the derivative of the input parameter as it will be used in `expression`.

For constitutive laws with more than one dimension, a string for each `epsilon_prime` and one for each `expression` are expected.

2.9.25 symbolic viscoelastic

The syntax is

```
<specific_const_law> ::= symbolic viscoelastic ,  
    epsilon , " <epsilon> " [ , ... ] ,  
    epsilon_prime , " <epsilon_prime> " [ , ... ] ,  
    expression , " <expression> " [ , ... ]
```

where `epsilon` and `epsilon_prime` are the symbols describing the input parameter and its derivative as they will be used in `expression`.

For constitutive laws with more than one dimension, a string for each `epsilon`, one for each `epsilon_prime`, and one for each `expression` are expected.

2.9.26 ann elastic

The implementation of the family of `ann` constitutive laws is based on an Artificial Neural Network library that is embedded in MBDyn.

The syntax is

```
<specific_const_law> ::= ann elastic ,  
    " <file_name> "
```

where the file `file_name` contains the parameters of the network. What is mainly significant to users is the need to scale inputs and outputs to match the amplitudes and possibly the offsets used in the training. For this purpose, the last two rows of the input files contain coefficients

```
b1 b0  
a1 a0
```

which are used to scale the input u and the output y according to the transformation

$$\bar{u} = b_1 u + b_0 \tag{2.58}$$

$$\bar{y} = a_1 y + a_0 \tag{2.59}$$

so that the actual output is

$$y = \frac{f(b_1 u + b_0) - a_0}{a_1} \tag{2.60}$$

2.9.27 ann viscoelastic

The syntax is

```
<specific_const_law> ::= ann viscoelastic ,
    " <file_name> "
```

where the file `file_name` contains the parameters of the network.

2.9.28 nlsf viscoelastic

This constitutive law was sponsored by Hutchinson CdR.

For an n -dimensional constitutive law, it implements the formula

$$f_i = \sum_{j=1,n} k'_{0ij} \varepsilon_j + f'_i(\varepsilon_i) + \sum_{j=1,n} k''_{0ij} \dot{\varepsilon}_j + f''_i(\dot{\varepsilon}_i) \quad (2.61)$$

where f'_i and f''_i are arbitrary instances of the `scalar functions` primitive, which includes piecewise linear and spline regularization of given data, while k'_{0ij} and k''_{0ij} are the constant coefficients of a linear viscoelastic model, that account for the cross-couplings between the stresses and the strains and strain rates.

Elastic and viscous variants are defined. They differ from the viscoelastic one by only allowing the specific fraction of the input data.

Syntax. The syntax is

```
<specific_const_law> ::= nlsf viscoelastic ,
    (DerivativeOfEntity) <kappa_0'> ,
    { null | (ScalarFunction) <diag_force'> }
    [ , ... ] ,
    { (DerivativeOfEntity) <kappa_0''> | proportional , <coef> } ,
    { null | (ScalarFunction) <diag_force''> }
    [ , ... ]

<specific_const_law> ::= nlsf elastic ,
    (DerivativeOfEntity) <kappa_0'> ,
    { null | (ScalarFunction) <diag_force'> }
    [ , ... ]

<specific_const_law> ::= nlsf viscous ,
    (DerivativeOfEntity) <kappa_0''> ,
    { null | (ScalarFunction) <diag_force''> }
    [ , ... ]
```

The terms `kappa_0'` and `kappa_0''` are indicated as `DerivativeOfEntity` because in the generic case they are the result of a differential operator that computes the derivative of the force vector as function of the strain or strain rate vectors. So, for a 3D constitutive law, the force, the strain and the strain rate are 3×1 vectors, while `kappa_0'` and `kappa_0''` are 3×3 matrices. Matrix input, in MBDyn, requires to write the whole set of coefficients, row-wise; however, it can be described in many synthetic manners if matrices have some special properties, as typical linear constitutive laws do.

For example, symmetric, diagonal and empty matrices have specific short forms; in the 3×3 case they are

```

# generic 3x3 matrix:
10., -2., -2.,
-2., 20., -8.,
-2., -8., 20.

# symmetric 3x3 matrix (same as above matrix):
sym,
10., -2., -2.,
    20., -8.,
        20.

# diagonal 3x3 matrix (diagonal of above matrix):
diag,
10., 20., 20.

# empty 3x3 matrix:
null

```

The terms `diag_force'` and `diag_force''` refer to the names of the scalar functions that are defined for each component of the force vector. There must be as many terms as the dimensions of the constitutive law. These terms can either be `null`, if no non-linear contribution is defined for that component, or contain a string, enclosed in double quotes, which must refer to an already defined scalar function. The same scalar function can be used multiple times.

Example. A 1D and a 3D constitutive law using the same exponential scalar function:

```

# define a scalar function
scalar function: "exponential", exp, coefficient, -2., 5.;

# define a 1D constitutive law
constitutive law: 1000, 1,
  nlsf viscoelastic,
    10.,
      null,          # stiffness is just linear
    0.,
      "exponential"; # damping is exponential

# define a 3D constitutive law
constitutive law: 3000, 3,
  nlsf viscoelastic,
    sym,
      10., -2., -2.,
        10., -2.,
          10.,
      null,          # stiffness is just linear
      null,          # stiffness is just linear
      null,          # stiffness is just linear
    null,
      "exponential", # damping is exponential...
      null,          # ...but in direction 1 only!
      null;

```

The first constitutive law corresponds to

$$f = 10.0 \cdot \varepsilon + 5.0 \cdot e^{-2\varepsilon} \quad (2.62)$$

while the second one corresponds to

$$\mathbf{f} = \begin{bmatrix} 10.0 & -2.0 & -2.0 \\ -2.0 & 10.0 & -2.0 \\ -2.0 & -2.0 & 10.0 \end{bmatrix} \begin{Bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \end{Bmatrix} + \begin{Bmatrix} 5.0 \cdot e^{-2\varepsilon_1} \\ 0.0 \\ 0.0 \end{Bmatrix} \quad (2.63)$$

2.9.29 nlp viscoelastic

This constitutive law was sponsored by Hutchinson CdR.

This constitutive law consists in the linear combination of nonlinear elastic and viscous effects whose coefficients are arbitrary nonlinear functions of the strain ε :

$$f_i = \sum_{j=1,3} (k'_{0ij} + \delta_{ij} k'_i(\varepsilon_i)) \varepsilon_j + \sum_{j=1,3} (k''_{0ij} + \delta_{ij} k''_i(\varepsilon_i)) \dot{\varepsilon}_j, \quad (2.64)$$

where δ_{ij} is Kronecker's operator; k'_i and k''_i are implemented as arbitrary **scalar functions**, which includes piecewise linear and spline regularization of given data.

Elastic and viscous variants are defined. They differ from the viscoelastic one by only allowing the specific fraction of the input data.

Syntax. The syntax is

```
<specific_const_law> ::= nlp viscoelastic ,
  (DerivativeOfEntity) <kappa_0'> ,
  { null | (ScalarFunction) <diag_stiffness> }
  [ , ... ] ,
  { (DerivativeOfEntity) <kappa_0''> | proportional , <coef> } ,
  { null | (ScalarFunction) <diag_damping> }
  [ , ... ]

<specific_const_law> ::= nlp elastic ,
  (DerivativeOfEntity) <kappa_0'> ,
  { null | (ScalarFunction) <diag_stiffness> }
  [ , ... ]

<specific_const_law> ::= nlp viscous ,
  (DerivativeOfEntity) <kappa_0''> ,
  { null | (ScalarFunction) <diag_damping> }
  [ , ... ]
```

Example. A 1D and a 3D constitutive law using the same exponential scalar function:

```
# define a scalar function
scalar function: "exponential", exp, coefficient, -2., 5.;

# define a 1D constitutive law
```

```

constitutive law: 1000, 1,
  nlp viscoelastic,
  10.,
    null,          # stiffness is just linear
  0.,
    "exponential"; # damping slope is exp.

# define a 3D constitutive law
constitutive law: 3000, 3,
  nlp viscoelastic,
  sym,
  10., -2., -2.,
    10., -2.,
    10.,
  null,          # stiffness is just linear
  null,          # stiffness is just linear
  null,          # stiffness is just linear
  null,
  "exponential", # damping slope is exp...
  null,          # ...but in direction 1 only!
  null;

```

The first constitutive law corresponds to

$$f = 10.0 \cdot \varepsilon + 5.0 \cdot e^{-2\varepsilon} \cdot \dot{\varepsilon} \quad (2.65)$$

while the second one corresponds to

$$\mathbf{f} = \begin{bmatrix} 10.0 & -2.0 & -2.0 \\ -2.0 & 10.0 & -2.0 \\ -2.0 & -2.0 & 10.0 \end{bmatrix} \begin{Bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \end{Bmatrix} + \begin{bmatrix} 5.0 \cdot e^{-2\varepsilon_1} & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 \end{bmatrix} \begin{Bmatrix} \dot{\varepsilon}_1 \\ \dot{\varepsilon}_2 \\ \dot{\varepsilon}_3 \end{Bmatrix} \quad (2.66)$$

Note: although the syntax of this constitutive law and that of the `nlsf viscoelastic` one (Section 2.9.28) is essentially identical, their behavior may be quite different, as indicated by Equations (2.61) and (2.64).

2.9.30 array

This constitutive law wrapper linearly combines the output of multiple constitutive laws. The syntax is

```

<specific_const_law> ::= array ,
  <number> ,
  <wrapped_const_law> [ , ... ]

```

The type of the resulting constitutive law is computed by combining that of the underlying constitutive laws. For example, if an elastic and a viscous constitutive law are combined, the type of the resulting array is viscoelastic. If `number` is 1, the underlying constitutive law is instantiated, and the array is ignored. Each underlying constitutive law can have a specific value of prestress and prestrain, as allowed by its own definition. No overall prestress nor prestrain is allowed by the `array` wrapper.

Example.

```

# the constitutive law
array, 2,
  linear elastic, 1000.,
  linear viscous, 10.

# is equivalent to
linear viscoelastic, 1000., 10.

```

2.9.31 bistop

This wrapper applies the logic of the bistop, that is activate or deactivate the constitutive property based on distinct conditions, with memory, to a generic underlying constitutive law. The syntax is

```

<specific_const_law> ::= bistop ,
  [ initial status , { inactive | active | (bool) <status> } , ]
  (DriveCaller) <activating_condition> ,
  (DriveCaller) <deactivating_condition> ,
  (ConstitutiveLaw<Entity, DerivativeOfEntity>) <wrapped_const_law>

```

2.9.32 invariant angular

This is not a constitutive law, but rather a wrapper for constitutive laws used within the “attached” variant of the **deformable hinge** joint. As such, it can only be used with the 3D dimensionality. It basically allows to refer an ancillary 3D constitutive law to an orientation that is intermediate with respect to the orientations of the two connected nodes.

```

<specific_const_law> ::= invariant angular ,
  <xi> , <wrapped_const_law>

```

xi is the fraction of relative orientation the ancillary constitutive law will be referred to, with respect to the first node of the joint. Its value should be comprised between 0 (attached to the first node) and 1 (attached to the second node); a value of 1/2 yields the **invariant deformable hinge** joint, but any value is allowed, not limited to within the [0, 1] range.

Note: the contribution to the Jacobian matrix will be incomplete when the underlying constitutive law has a viscous contribution. This may slow down convergence, or even prevent it. If this is an issue, the **invariant deformable hinge** should be used, since it does not suffer from this limitation, although it only allows **xi** = 1/2.

2.10 Hydraulic fluid

Hydraulic fluid data defines the constitutive properties of hydraulic fluids, which are generally required by hydraulic elements. Hydraulic fluid data can be defined in two ways, according to the BNF:

```

<hydraulic_fluid> ::=
  { <fluid_type> , <fluid_properties>
    | reference , <label> }

```

The latter references a previously defined hydraulic fluid dataset, described in Section 2.3.4. The available types **fluid_type**, with the related **fluid_properties**, are:

2.10.1 Incompressible

```
<fluid_type> ::= incompressible

<fluid_properties> ::=
  <property> [ , ... ]

<property> ::=
  { density , <density>
    | viscosity , <viscosity>
    | pressure , <pressure>
    | temperature , <temperature> }
```

2.10.2 Linearly compressible

```
<fluid_type> ::= linear compressible

<fluid_properties> ::=
  <property> [ , ... ]

<property> ::=
  { density , <ref_density> ,
    { sound celerity , <sound_celerity> | <beta> } , <ref_pressure>
    | viscosity , <viscosity>
    | temperature , <temperature> }
```

If the keyword `sound celerity` is used, the bulk modulus `beta` is computed as

$$\text{beta} = \text{ref_density} \cdot \text{sound_celerity}^2 \quad (2.67)$$

2.10.3 Linearly compressible, with thermal dependency

```
<fluid_type> ::= linear thermal compressible

<fluid_properties> ::=
  <property> [ , ... ]

<property> ::=
  { density , <ref_density> ,
    { sound celerity , <sound_celerity> | <beta> } , <ref_pressure>
    <alpha> , <ref_temperature>
    | viscosity , <viscosity> }
```

2.10.4 Super (linearly compressible, with thermal dependency)

```
<fluid_type> ::= super

<fluid_properties> ::=
  <property> [ , ... ]

<property> ::=
```

```

{ density , <ref_density> ,
  { sound_celerity , <sound_celerity> | <beta> } , <ref_pressure>
| viscosity , <viscosity>
| temperature , <temperature> }

```

according to equation

$$\begin{aligned}
\rho_- &= \rho_0 + \text{ref_density} \cdot \frac{1}{2}(1 + \tanh(a(p - p_0))) \\
\rho &= \rho_- & p < p_0 \\
\rho &= \rho_- + \frac{p - p_0}{\text{beta}} & p > p_0
\end{aligned} \tag{2.68}$$

Note: highly experimental; currently, ρ_0 , a and p_0 are hardcoded in SI units; the syntax will change

2.10.5 Exponential compressible fluid, with saturation

```

<fluid_type> ::= exponential

<fluid_properties> ::=
  [ <property> [ , ... ] , ] <p_sat>

<property> ::=
  { density , <ref_density> ,
    { sound_celerity , <sound_celerity> | <beta> } , <ref_pressure>
  | viscosity , <viscosity>
  | temperature , <temperature> }

```

where `p_sat` is the saturation pressure, according to equation

$$\begin{aligned}
\rho &= \text{ref_density} \cdot e^{\frac{p - \text{ref_pressure}}{\text{beta}}} & p > \text{p_sat} \\
\rho &= \text{ref_density} \cdot e^{1000 \frac{p - \text{ref_pressure}}{\text{beta}}} & p < \text{p_sat}
\end{aligned} \tag{2.69}$$

Note: this fluid constitutive law is loosely inspired by AMESim's simply corrected compressible fluid.

2.11 Authentication Methods

Some authentication methods are defined and made available to specific program modules; they are used to authenticate before accessing some resources of the program while it is running. The syntax is:

```

<authentication_method> ::= <method> [ , <specific_data> ]

```

Authentication methods in general expect some authentication tokens to be input. Usually a user name and a password are required.

2.11.1 Note on security and confidentiality

No encryption is used in communications, unless provided by the underlying mechanism (e.g. some SASL mechs), so the authentication methods are very rough and should be considered as insecure. TLS may be considered in the future. If confidentiality is required, SASL with at least DIGEST-MD5 is strongly recommended; otherwise no authentication should be used. As an alternative, a SSH tunnel may be established between the client and the server machine, and simple authentication can be used. Otherwise, if the user has direct access to the server where the computation is being run, sockets with local namespace can be used, and security can be enforced by means of the access privileges of the socket file. Since some of the UN*X systems do not honor socket permissions, a portable way to exploit filesystem access permissions is to put the socket in a dedicated directory, and use the permissions of the directory to control access to the socket.

Available methods are:

2.11.2 No authentication

```
<authentication_method> ::= no auth
```

2.11.3 Password

```
<authentication_method> ::= password

<specific_data> ::=
  user , " <user_name> " ,
  credentials , { prompt | " <user_cred> " }
  [ , salt_format , <salt_format> ]
```

In case the keyword `prompt` is given as `credentials`, the user is prompted for a password; `user_cred` is used otherwise. The optional parameter `salt_format` allows to specify different formats for the `salt`, for those `crypt(3)` extensions that support more sophisticated encryption mechanisms (e.g. MD5). See `crypt(3)` for details. If the credentials are preceded by the string `{CRYPT}`, they are assumed to be already encrypted, and the remaining portion is used.

2.11.4 PAM (Pluggable Authentication Modules)

```
<authentication_method> ::= pam

<specific_data> ::=
  [ user , " <user_name> " ]
```

The *Linux-PAM* Pluggable Authentication Modules can be used to authenticate a user. If no user name is provided, the effective user id, as provided by the `geteuid(2)` system function, is used to retrieve the username of the owner of `mbdyn` process. the `user` must be valid. The authentication is performed through a system-dependent `pam` configuration file. No checks on the validity of the account or on the permission of opening a session are made; account, session and password changes should be explicitly denied to `mbdyn` to avoid possible security breaks (see the following example). The interested reader should consult the documentation that comes with the package, try for instance <http://www.kernel.org/pub/linux/libs/pam/> for details.

An example is provided with the package, in `/etc/pam.d/mbdyn`:

```

### use either of the following:
auth      required    /lib/security/pam_unix_auth.so
# auth    required    /lib/security/pam_pwdb.so
#
### no account, session or password allowed
account   required    /lib/security/pam_deny.so
session   required    /lib/security/pam_deny.so
password  required    /lib/security/pam_deny.so

```

which allows authentication by using standard Un*x or `libpwdb` based authentication.

2.11.5 SASL (Simple Authentication and Security Layer)

```

<authentication_method> ::= sasl

<specific_data> ::= [ <parameter> [ , ... ] ]

<parameter> ::=
    { user , " <user_name> "
      | mechanism , " <preferred_mechanism> " }

```

This is the preferred authentication method because it is mechanism-independent, it can be reasonably secure and automatically selects the most appropriate mechanism available on both the client and the server machine. It requires Cyrus SASL 2 (See <http://asg.web.cmu.edu/sasl/> for details, and follow the documentation to obtain a working setup).

2.12 Miscellaneous

Finally there are some miscellaneous points:

- (UN*X systems) Environment variables whose name starts with MBDYN may be defined and passed to an execution of the mbdyn command. The following are recognized at present:
 1. `MBDYNVARS=<expr_list>` where `expr_list` is a series of mathematical expressions separated by semicolons. They are parsed and evaluated; if variables are declared, they are added to the symbol table to be used during the whole execution of the program.
 2. `MBDYN_<type>_<name>=<value>`, where `type` is a legal mbdyn type (`integer`, `real` or `string`; see Table 2.3 for details), `name` is a legal symbol name and `value` is a legal expression.

Example: sample code to execute MBDyn from within a python script, passing it the time step in the `real` variable DT:

```

import os;
dt = 1.e-3;
os.environ['MBDYNVARS'] = 'real DT = ' + str(dt) + ';';
os.system('mbdyn -f input.mbd -o output > output.txt 2>&1 &');
# use 'dt' in the python peer's analysis as appropriate...

```

- Newlines and indentations are not meaningful. But good indentation habits can lead to better and more readable input files.

- Everything that follows the character ‘#’ is considered a remark, and is discarded until the end of the line. This can occur everywhere in the file, even inside a math expression (if any problems occur, please let me know, because chances are it is a bug!)
- A new style for comments has been introduced, resembling the C programming language style: everything comprised between the marks /* and */ is regarded as a remark:

```
/*
 * useful comments make input files readable!
 */
```

This can happen (almost) everywhere in the text except in the middle of a keyword.

- (UN*X systems) Whenever a file name contains a portion of the form \$VARNAME or \${VARNAME}, appropriate expansion from environment is performed; VARNAME is

```
VARNAME ::= [_a-zA-Z][_a-zA-Z0-9]*
```

namely, it must begin with a letter or an underscore, and can be made of underscores, letters and digits.

- (UN*X systems) Whenever a file name is required, if the string containing the file name starts with '\$', shell environment variable expansion occurs.

Example: expand the local socket’s name in `external` forces

```
# ...
path, "$MBSOCK",
# ...
```

Example: sample code to execute MBDyn from within a python script, passing it the name of the local socket to be used for an `external` force:

```
import os;
import tempfile;
tmpdir = tempfile.mkdtemp('', '.mbdyn_');
path = tmpdir + '/mbdyn.sock';
os.environ['MBSOCK'] = path;
os.system('mbdyn -f input.mbd -o output > output.txt 2>&1 &');
# use 'path' as the name of the socket as shown in the previous example
# ...
# when done...
os.rmdir(tmpdir);
```

This code:

- creates a uniquely named temporary directory in the default temporary directory (e.g. /tmp) only accessible to the owner of the process, according to /tmp/.mbdyn_XXXXXX, where XXXXXX are six random chars (see `mkdtemp(3)` for details);
- appends the name /mbdyn.sock to form the full path of the socket;
- sets the full path of the socket in the environment under the name MBSOCK;

- executes MBDyn; its input file is supposed to use "\$MBSOCK" as path name of a local socket.

This procedure guarantees that a co-simulation performed between MBDyn and the peer Python script uses a unique, protected socket and thus does not interfere with other co-simulations running on the same system.

- (UNIX systems) Whenever a file name is required, the shell-like syntax for home directories (i.e. ~/filename or ~user/filename is automatically resolved if legal [user and] filename values are inserted. Home expansion occurs after environment variable expansion (see above).
- The **license** and the **warranty** statements respectively show on the standard output the license and the warranty statement under which the code is released. They do not affect the simulation.

Chapter 3

Data

The `data` section is read directly by the driver program. It is included between the cards:

```
begin : data ;  
    # ...  
end : data ;
```

3.1 Problem

The only active card is the `problem` card.

The syntax is

```
<card> ::= problem : <problem_name> ;  
  
<problem_name> ::= { initial value | inverse dynamics } ;
```

For long time the only supported problem type was the `initial value` problem. As a consequence, if no `problem` card is present, an initial value problem is assumed. An experimental `inverse dynamics` problem is supported as well.

Chapter 4

Problems

This section is used to insert all the data related to the problem that one needs MBDyn to solve in the simulation. The section data is included between the cards:

```
begin : <problem_name> ;  
# ...  
end : <problem_name> ;
```

Implemented problems are:

- **initial value**, the time integration of mechanical and multidisciplinary problems formulated as Differential-Algebraic Equations (DAE). It can be downgraded to the solution of static and kinematic problems, by selecting purely static or kinematic contributions to the governing equations, thus giving time the role of an ordinal parameter.
- **inverse dynamics**, the computation of the generalized forces required to make a generic system perform a given trajectory. This problem is currently under development, so it is not discussed in detail.

4.1 Initial-Value Problem

At present, the main problem is **initial value**, which solves initial value problems by means of generic integration schemes that can be casted in a broad family of multistep and, experimentally, Implicit Runge-Kutta-like schemes [5].

The syntax of the module is:

```
begin : initial value ;  
# ...  
end : initial value ;
```

At present, there are a number of cards that can be grouped as follows, based on the integration phase they refer to.

4.1.1 General Data

those data that refer to the main integration phase or the simulation as a whole. They are:

Initial Time

```
<card> ::= initial time : <time> ;
```

Final Time

```
<card> ::= final time : { forever | <time> } ;
```

Strategy

```
<card> ::= strategy : <strategy_type> [ , <strategy_data> ] ;
```

The available strategies are:

- **no change**

```
<strategy_type> ::= no change
```

the time step is never changed. However, this prevents simulation termination when the maximum number of iterations is reached. As a consequence, its use only makes sense when repeating a time step results in some modification of the analysis related to some user-defined function.

- **factor**

```
<strategy_type> ::= factor
```

```
<strategy_data> ::=  
  <reduction_factor> ,  
  <steps_before_reduction> ,  
  <raise_factor> ,  
  <steps_before_raise> ,  
  <min_iterations>  
  [ , <max_iterations> ]
```

the time step is reduced or raised of the proper factor not before a minimum number of time steps; it is reduced if more than `max_iterations` are performed at a time step, or if the current step does not converge; it is raised if less than `min_iterations` are performed at a time step; `max_iterations` defaults to the global `max_iterations` simulation value; however, it is better to set it to a lower value, leaving some spare iterations before bailing out the simulation; the simulation bails out if two consecutive solution steps are performed without converging.

- **change**

```
<strategy_type> ::= change
```

```
<strategy_data> ::= (DriveCaller) <time_step_pattern>
```

the time step is changed according to the `time_step_pattern` pattern. If the time step returned by the drive caller does not change when a step needs to be repeated, the analysis is interrupted.

Note: the `change` strategy is not normally intended to provide adaptive time step, it rather allows to prescribe a given variable time step pattern based on the rule defined by the `time_step_pattern DriveCaller`.

In any case, step change only occurs after the first step, which is performed using the `time_step` value provided with the `time step` statement.

Min Time Step

```
<card> ::= min time step : <min_time_step> ;
```

Defines the minimum value allowed for the time step. This statement is only meaningful when variable time step is used. If the time step change strategy tries to use a time step smaller than `min_time_step`, the simulation aborts.

Max Time Step

```
<card> ::= max time step : { <max_time_step> | unlimited } ;
```

Defines the maximum value allowed for the time step. This statement is only meaningful when variable time step is used. If the time step change strategy tries to use a time step larger than `max_time_step`, the value `max_time_step` is used instead.

Time Step

```
<card> ::= time step : <time_step> ;
```

The initial time step. This value is used throughout the simulation unless some variable time step strategy is defined using the `strategy` statement.

Tolerance

```
<card> ::= tolerance : { null | <residual_tolerance> }  
          [ , test , { none | norm | minmax } [ , scale ] ]  
          [ , { null | <solution_tolerance> }  
            [ , test , { none | norm | minmax } ] ] ;
```

The only mandatory value is `residual_tolerance`, the tolerance used for the residual test; the keyword `null` disables the residual testing, disabling the computation of the test on the residual. The `test` mechanism is used to select what kind of test must be performed; currently, only `norm` (the default) and `minmax` are supported. The special value `none` means that the test is not actually computed; it is the default when the test is disabled by setting the tolerance to zero, or by using the keyword `null`; however, it can be restored to any mechanism for output purposes. The optional parameter `scale` is used to enable the scaling of the residual before performing the test; default scale factors can be set for each type of degree of freedom owner, and some entities allow individual scaling; by default, no scaling takes place. A tolerance `solution_tolerance` to test the solution (the difference between the states at two iterations) is allowed, and also in this case a test mechanism can be chosen; by default, no test on the solution convergence is done. Currently, no scaling is allowed in the solution test.

Example.

```
# residual test by means of the norm  
tolerance: 1.e-6;  
# residual test with minmax method  
tolerance: 1.e-6, test, minmax;  
# residual test with norm method and scaling  
tolerance: 1.e-6, test, norm, scale;  
# solution test  
tolerance: null, 1.e-9;
```

```

# residual and solution test
# (the first that succeeds breaks the loop)
tolerance: 1.e-6, 1.e-9;
# residual test with computation of solution norm
tolerance: 1.e-6, null, test, norm;

```

Max Iterations

```

<card> ::= max iterations : <max_iterations> [ , at most ] ;

```

Error out after `max_iterations` without passing the convergence test.

If the optional keywords `at most` are given, the step is considered successfully performed after `max_iterations` regardless of passing the convergence test, as soon as the error reduced since the first iteration. For example, when using

```

max iterations: 5, at most;

```

the convergence pattern

```

...
Iteration(0) 1.16374e+06 J
      SolErr 0
Iteration(1) 8330.76 J
      SolErr 0
Iteration(2) 43768.8 J
      SolErr 0
Iteration(3) 2981.01 J
      SolErr 0
Iteration(4) 118839 J
      SolErr 0
Iteration(5) 8654.02 J
...

```

succeeds, since $8654.02 < 1.16374e+06$, while

```

...
Iteration(0) 1.66901e+06 J
      SolErr 0
Iteration(1) 2.06327e+09 J
      SolErr 0
Iteration(2) 2.64337e+10 J
      SolErr 0
Iteration(3) 1.2131e+12 J
      SolErr 0
Iteration(4) 7.08449e+12 J
      SolErr 0
Iteration(5) 2.18828e+15 J

```

Max iterations number 5 has been reached during Step=6, Time=0.06; \ TimeStep=0.01 cannot be reduced further; aborting...

An error occurred during the execution of MBDyn; aborting...

fails since $2.18828e+15 > 1.66901e+06$.

This option should be used with extreme care.

Modify Residual Test

```
<card> ::= modify residual test ;
```

modify the residual test taking in account the rate of change of the status.

Method

```
<card> ::= method : <method_data> ;
```

there are four multistep methods at present.

The first is Crank-Nicolson:

```
<method_data> ::= crank nicolson
```

The second and the third are original methods; the former is discussed in [5], see also

```
http://www.aero.polimi.it/mbdyn/publications.html
```

Those methods are unconditionally stable and can be tuned to give the desired algorithmic dissipation by setting the value of the asymptotic spectral radius. The radius can be set independently for the differential and the algebraic variables, and a driver is used, to allow parameter dependent radius variation.

```
<method_data> ::= { ms | hope } ,  
  (DriverCaller) <differential_radius>  
  [ , (DriverCaller) <algebraic_radius> ]
```

The first method proved to be more accurate at high values of asymptotic radius (low dissipation), while the second proved to be more accurate at low values of the radius (high dissipation). They look nearly equivalent at radii close to 0.4, with the former giving the best compromise between algorithmic dissipation and accuracy at about 0.6. The `algebraic_radius` can be omitted, defaulting to the same as the differential one. It is unclear whether a different spectral radius can help in increasing accuracy or dissipation of the algebraic unknowns.

The fourth method is experimental. It is a third-order, two stage unconditionally stable method, which can be tuned to give the desired algorithmic dissipation by setting the value of the asymptotic spectral radius, which should not be too close to zero. Currently it is not possible to independently set the radius for the differential and the algebraic variables.

```
<method_data> ::= third order ,  
  { ad hoc | (DriverCaller) <differential_radius> }
```

When the keyword `ad hoc` is used, a method equivalent to a Runge-Kutta Radau IIA results, with zero asymptotic spectral radius; otherwise, an original third-order integration method is used, with tunable algorithmic dissipation.

The multistep method, when the asymptotic radius is zero, degenerates in the Backward Differentiation Formulas of order two. A shortcut to this case is provided as

```
<method_data> ::= bdf [ , order , <order> ]
```

The keyword `order` can be used to indicate a specific `order` of the BDF formulas; only first order (implicit Euler) and second order formulas are currently implemented, and the default is the second order formula, which is the most useful. The first order formula may be of help for very specific problems. It can also be selected using the shortcut

```
<method_data> ::= implicit euler
```

Nonlinear Solver

The nonlinear solver solves a nonlinear problem $F(x) = 0$. The syntax is

```
<card> ::= nonlinear solver : <nonlinear_solver_data> ;
```

Two nonlinear solvers are currently available: a Newton-Raphson iterator and a matrix-free solver.

The first nonlinear solver, a modified Newton-Raphson, is well tested.

```
<nonlinear_solver_data> ::= newton raphson :  
  { true | modified , <iterations>  
    [ , keep jacobian matrix ]  
    [ , honor element requests ] } ;
```

if **modified**, the number of **iterations** the same Jacobian matrix will be reused, and thus factored only once, is expected. If the option **keep jacobian matrix** is selected, the Jacobian matrix is preserved for the desired **iterations** even across time steps. By default, the Jacobian matrix is recomputed at the beginning of each time step. If the option **honor element requests** is selected, the factorization is updated also when an element changes the structure of its equations. The default behavior is to ignore such requests¹.

The second nonlinear solver, based on a matrix-free algorithm, is experimental.

```
<nonlinear_solver_data> ::= matrix free , { bicgstab | gmres }  
  [ , tolerance , <tolerance> ]  
  [ , steps , <steps> ]  
  [ , tau , <tau> ]  
  [ , eta , <eta> ]  
  [ , preconditioner ,  
    { full jacobian matrix [ , steps , <steps> ]  
      [ , honor element requests ] } ] ;
```

where **tolerance** is the iterative linear solver tolerance; **steps** is the maximum number of linear solver iterations; **tau** is a measure of the configuration perturbation used to obtain a matrix-free approximation of the product of the Jacobian matrix times the solution vector; it is seldom necessary to change the default value of this parameter, mainly for very “stiff” problems. By setting **eta** the convergence requirements are changed; further reference for this parameter may be found in [6]; do not change it unless you know what you are doing. The value **steps** of the **steps** keyword of the **preconditioner** sub-block sets the maximum number of iterations that can be performed without requiring an update of the preconditioner; this parameter can heavily influence the performance of the matrix-free solver. If the option **honor element requests** is selected, the preconditioner is updated also when an element changes the structure of its equations. The default behavior is to ignore such requests².

Eigenanalysis

```
<card> ::= eigenanalysis :  
  { <when> | list , <num_times> , <when> [ , ... ] }  
  [ , { output [ full ] matrices  
        | output sparse matrices  
        | output eigenvectors  
        | output geometry
```

¹ At present, only few elements that can change the structure of the equations, or at least radically change the Jacobian matrix, actually issue this request.

²See note above

```

    | parameter , <param>
    | lower frequency limit , <lower>
    | upper frequency limit , <upper>
    | <method> }
    [ , ... ] ]
<method> ::=
    { use lapack [ , balance , { no | scale | permute | all } ]
      | use arpack , <nev> , <ncv> , <tol>
      | use jdqz , <nev> , <ncv> , <tol> }

```

Performs the direct eigenanalysis of the problem. This functionality is experimental; see [7] for theoretical aspects. Direct eigenanalysis based on the matrices of the system only makes sense when the system is in a steady configuration, so the user needs to ensure this configuration has been reached. Moreover, not all elements currently contribute to the Jacobian matrix of the system, so YMMV. In case of rotating systems, a steady configuration could be reached when the model is expressed in a relative reference frame, using the **rigid body kinematics** card (see Section 5.2.15).

- **when** indicates at what time the eigenanalysis is performed. The analysis is performed at the first time step for which $\text{Time} \geq \text{when}$. If the keyword **list** is given, **num_times** values of **when** are expected. At most **num_times** eigenanalyses are performed; the output file name is composed using the position of the corresponding **when** in the array of times; if **when** < **initial_time** the first eigenanalysis is performed before the “derivatives” phase; if **when** = **initial_time** the first eigenanalysis is performed after the “derivatives” phase and before the first regular step;
- **param** is the value of the coefficient that is used to project the problem in the discrete time domain; it represents a singularity point in the transformation from the discrete to the continuous eigenvalues, so it should not coincide with an eigenvalue of the problem; by default it is 1.0;
- **output full matrices** causes the output of the (full) problem matrices in a .m file, intended to be loaded by tools like Octave/Matlab;
- **output sparse matrices** causes the output of the problem matrices in sparse form in a .m file, intended to be loaded by tools like Octave/Matlab;
- **output eigenvectors** causes the output of the eigenvalues and of the eigenvectors in a .m file, intended to be loaded by tools like Octave/Matlab;
- **output geometry** causes the output of the reference configuration of structural nodes, if any, and of the indexes that need to be used in order to extract the perturbation of their position and orientation from the eigenvectors;
- **use lapack** performs the eigenanalysis using LAPACK’s **dggev()** routine (generalized unsymmetric eigenanalysis using the QZ decomposition);
- **balance** performs matrix permutation (**permute**), scaling (**scale**), both (**all**) using LAPACK’s **dggev()** routine, or none (**no**) using **dggev()**; by default, matrix permutation is performed, using **dggev()**;
- **use arpack** performs the eigenanalysis using ARPACK’s **dnaupd()** and **dneupd()** routines (canonical unsymmetric eigenanalysis, using the Implicitly Restarted Arnoldi Method, IRAM). It requires the additional parameters:
 - **nev**, the number of the desired eigenvalues;

- `ncv`, the number of the Arnoldi vectors;
- `tol`, the tolerance (positive; zero means machine error).
- `use jdqz` performs the eigenanalysis using JDQZ (Jacobi-Davidson QZ decomposition). It requires the same parameters of ARPACK; actually, it allows a lot of tuning, but its input is not formalized yet (experimental).
- `lower frequency limit` only outputs those eigenvalues and eigenvectors that are complex and whose imaginary part is larger than `lower` Hz;
- `upper frequency limit` only outputs those eigenvalues and eigenvectors that are complex and whose imaginary part is smaller than `upper` Hz;

The `eigenanalysis` functionality is experimental, so the syntax could need adjustments and change across versions.

Note: in practical cases, the use of lapack with `balance` appeared to be quite sensitive to the problem data. In some cases, the use of `scale` may introduce errors in eigenvalues. The use of `permute` is usually beneficial: it allows to directly single out some of those spurious eigenvalues that correspond to constraint equations.

Note: in practical cases, the usefulness of arpack appeared to be questionable, especially in those cases where a significant amount of constraints is present. It may be beneficial in those cases where the number of deformable components is much larger than the number of constraint equations.

Output A brief summary of the eigenvalues is output in the `.out` file. Most of the information is dumped in a `.m` file in a format that is suitable for execution with scientific software like Octave and Matlab. The execution of the file populates the environment of those software with a set of variables that contains the results of the eigenanalysis. According to the options used in the `eigenanalysis` card, the `.m` file may contain:

- `Aplus`: matrix $\mathbf{J}_{(+c)} = \mathbf{f}_{/\dot{\mathbf{x}}} + \text{dCoef} \cdot \mathbf{f}_{/\mathbf{x}}$, either in full or sparse format;
- `Aminus`: matrix $\mathbf{J}_{(-c)} = \mathbf{f}_{/\dot{\mathbf{x}}} - \text{dCoef} \cdot \mathbf{f}_{/\mathbf{x}}$, either in full or sparse format;
- `dCoef`: the coefficient used to build the matrices $\mathbf{J}_{(+c)}$ and $\mathbf{J}_{(-c)}$;
- `alpha`: an array of three columns containing $\alpha_r = \text{alpha}(:, 1)$, $\alpha_i = \text{alpha}(:, 2)$ and $\beta = \text{alpha}(:, 3)$, such that the k -th discrete-time eigenvalue Λ_k is

$$\Lambda_k = \frac{\alpha_r(k) + i\alpha_i(k)}{\beta(k)}. \quad (4.1)$$

Note that α_r , α_i and β can be all zero; only a fraction of the eigenvalues may be output; the corresponding continuous time domain eigenvalue λ_k is

$$\lambda_k = \frac{1}{\text{dCoef}} \frac{\Lambda_k - 1}{\Lambda_k + 1} = \frac{1}{\text{dCoef}} \frac{\alpha_r(k)^2 + \alpha_i(k)^2 - \beta(k)^2 + i2\alpha_i(k)\beta(k)}{(\alpha_r(k) + \beta(k))^2 + \alpha_i(k)^2}, \quad (4.2)$$

where `dCoef` is defined in the `.m` file;

- `VR`: the right eigenvectors, solution of the problem $(\mathbf{J}_{(+c)} + \Lambda\mathbf{J}_{(-c)}) \mathbf{V}_R = \mathbf{0}$; only the eigenvectors corresponding to the selected eigenvalues are output;
- `VL`: the left eigenvectors, solution of the problem $(\mathbf{J}_{(+c)} + \Lambda\mathbf{J}_{(-c)})^T \mathbf{V}_L = \mathbf{0}$; only the eigenvectors corresponding to the selected eigenvalues are output;

- **X0**: a vector containing the reference position and orientation of the structural nodes, if any (excluding any dummy node); the vector contains, for each node, the three components of the reference position and the three components of the Euler vector corresponding to the reference orientation; all data is referred to the absolute reference frame;
- **idx**: a vector containing the reference row index of each structural node in the eigenvectors;

The contents of this file may change across versions.

Example.

```
# to access the "x" component of each node of each eigenvector
octave:1> VR(idx + 1, :)
# to plot the "y" component of each node of each eigenvector,
# with respect to the reference "x" component of the position of the nodes
octave:2> plot(X0(1:6:end), VR(idx + 2, :))
```

Abort After

```
<card> ::= abort after :
  { input
    | assembly
    | derivatives
    | dummy steps } ;
```

mainly used to cleanly check models and simulations at various phases. When set to **input**, the simulation ends after input is completed; when set to **assembly**, the simulation ends after initial assembly, if required; when set to **derivatives**, the simulation ends after the computation of the derivatives, i.e. after the system is solved at the initial time to compute initial momentum and momenta moment derivatives and reaction forces; when set to **dummy steps**, the simulation ends after the dummy steps execution, if required. In any case, the output is generated with the system configured as resulting from the last computation phase, so this mode can be useful to check how the system is actually configured after phases that are not usually output.

Linear Solver

```
<card> ::= linear solver :
  { naive | umfpack | klu | y12 | lapack | superlu | taucs | watson }
  [ , { map | cc | dir } ]
  [ , { colamd | mmdata } ]
  [ , { mt | multithread } , <threads> ]
  [ , workspace size , <workspace_size> ]
  [ , pivot factor , <pivot_factor> ]
  [ , drop tolerance , <drop_tolerance> ]
  [ , block size , <block_size> ]
  [ , scale , { no | always | once } ]
  ;
```

Umfpack. The default, when available, is **umfpack**. Since it dynamically allocates memory as required, the **workspace size** parameter is ignored. The **umfpack** linear solver honors the **block size** keyword, which refers to an internal parameter. The default value (32) is usually good for most problems; however,

it has been observed that a lower value can slightly improve performances on small problems (e.g. a rigid rotorcraft model with some 180 degrees of freedom showed a 10% speedup with a block size of 4, while a rigid robot showed even higher speedups, of the order of 30%, with a block size between 4 and 16).

Naive. The `naive` solver is built-in, so it is always present. When `umfpack` is not available, it is used by default. The naive solver uses a full storage but keeps track of the non-zeros. This allows a very efficient solution of sparse matrices, with a $O(1)$ access cost to the coefficients. It ignores the `workspace size` parameter. See [8] for details.

Important note: the naive solver should always be used with some reordering option. The `colamd` option is always available (see discussion below), since it is built-in. Other options are experimental, and should be used with care. If this option is not used, the factorization may not be robust. If the `linear solver` statement is not used, and the `naive` solver is used by default (because no other solver with higher priority is present), **no reordering option is set by default**. In this case, it is safe to explicitly force the linear solver with a valid reordering option. For example

```
linear solver: naive, colamd;
```

Klu. The `klu` solver is provided by University of Florida's SparseSuite. It is very efficient (usually faster than `umfpack`), but it should be used with care, because in the current implementation, to exploit its efficiency, the symbolic factorization is reused across factorizations. Whenever the solver realizes that the symbolic factorization needs to be regenerated, because the structure of the matrix changed (e.g. when using `driven` elements, but not only in that case), the matrix is already invalidated, and should be regenerated. This is not handled correctly yet.

Y12. The `y12` linear solver is provided for historical reasons and should not be used for practical purposes. It requires static allocation of the workspace; the optimal size `workspace_size` is between 3 and 5 times the number of nonzero coefficients of the sparse matrix. By default, twice the size of the full matrix is allocated (which sounds a bit nonsense, but is conservative enough); then the linear solver automatically uses an optimal amount of memory based on the results of the previous factorization; this could result in a failure if the filling of the matrix changes dramatically between two factorizations.

Lapack. The `lapack` linear solver uses LAPACK's `dgetrf()` and `dgetrs()` functions with a dense storage, and thus ignores the `workspace size` parameter. It is fairly inefficient for most practical applications, when compared to sparse solvers, except when problems are very small (less than 60 equations) and very dense.

SuperLU. Finally, `superlu` is an experimental linear solver that is able to perform the factorization in a multi-threaded environment. If more than one CPU is available, it automatically uses all the CPUs; the keyword `mt` allows to enforce the desired number of `threads`. See also the `threads` keyword for more details on multi-threaded solution.

Other. Historically, MBDyn supported the `meschach` and `harwell` solvers. They are now deprecated and undocumented. Other linear solvers may be supported for experimental reasons (for example, `taucs`, `watson`); however, they are not documented as they are not intended for general use.

Miscellaneous keywords. The keywords `map` (Sparse Map), `cc` (Column Compressed) and `dir` (Direct Access) can be used to modify the way the sparse matrix is handled:

- `map` uses an array of binary trees to store the matrix coefficients; right before the factorization, the matrix is transformed in the form required by each linear solver, usually Column Compressed (also known as Harwell-Boeing format) or Index (as used by `y12`). The matrix is regenerated each time an assembly is required.
- `cc` uses the compact format resulting from the transformation required by the linear solver to store the coefficients in subsequent assemblies, thus saving the time required to transform the matrix. If the profile of the matrix changes, the compact form is invalidated, and the matrix is reset to `map` form. As a consequence, this matrix form allows faster access and reduced factorization cost, but in certain cases may require an additional cost.
- `dir` uses a full index table to grant direct access time to the actual location of each coefficient. In principle it should be a bit faster than `cc`, at the cost of some memory waste. However, in practical applications, the cost of these operations is a very small fraction of the factorization time. If the problem is quite large, the memory consumption might become unacceptable; if the problem is quite small, the access time of `cc` is quite limited ($\log_2(\langle \text{nz} \rangle)$, where `nz` is the number of non-zeroes in a column) so it is definitely comparable to that of `cc`.

Only `umfpack`, `klu`, `y12` and `superlu` linear solvers allow these settings.

The keyword `colamd` is honored by the `naive` and SuperLU linear solvers; it enables the column rearrangement that minimizes the sparsity of the factored matrix, as implemented in `libcolamd` (part of University of Florida's SparseSuite package). The keyword `mmdata` is honored by the scalar SuperLU linear solver only; it enables column rearrangement based on the MMD ordering of the symmetric matrix $A^T A$.

The `pivot_factor` is a real number, which in a heuristic sense can be regarded as the threshold for the ratio between two coefficients below which they are switched, so 0.0 means no pivoting, while 1.0 means switch as soon as the norm of the ratio is less than unity.

The keyword `drop tolerance` enables incomplete factorization. Elements of the factorization whose absolute value is below the `drop_tolerance` threshold are discarded. By default it is 0. Only `umfpack` allows this setting.

The keyword `block size` is specific to `umfpack`. See the related documentation for the meaning.

The keyword `scale` indicates whether the matrix and the right-hand side must be scaled before linear solution. By default no scaling takes place (`no`). Only `naive` allows this setting. It is experimental; it should only be used for development purposes.

Table 4.1 summarizes the properties that can be set for each type of linear solver; Table 4.2 summarizes the memory usage of the linear solvers; Table 4.3 summarizes how each linear solver deals with pivoting policies.

Solver

Deprecated; see `linear solver` (Section 4.1.1).

Threads

```
<card> ::= threads :  
      { auto | disable | [ { assembly | solver } , ] <threads> }
```

By default, if enabled at compile time, the assembly is performed in a multi-threaded environment if more than one CPU is available and the selected solver allows it (e.g. if it supports any form of compressed

Table 4.1: Linear solvers properties

<i>Solver</i>	<i>Map</i>	<i>Column Compr.</i>	<i>Dir</i>	<i>Multi- Thread</i>	<i>Worksp. Size</i>	<i>Pivot Factor</i>	<i>Block Size</i>
Naive				(√)		√	
Umfpack	√	√	√			√	√
KLU	√	√	√			√	
Y12m	√	√	√		√	√	
Lapack						√	
SuperLU	√	√	√	√		√	
TAUCS	?	?	?	?	?	?	√
Watson	√	√	√	√		√	

Table 4.2: Linear solvers memory usage

<i>Solver</i>	<i>Workspace</i>	<i>Memory Size</i>	<i>Block Size Allocation</i>
Naive		full	
Umfpack		dynamic	default=32
KLU		dynamic	
Y12m	default= $2 \times n^2$ $3nz \rightarrow 5nz$	static	
Lapack		full	
SuperLU		dynamic	
TAUCS		?	
Watson		dynamic	

Table 4.3: Linear solvers pivot handling

<i>Solver</i>	<i>Pivot</i>	<i>Default</i>	<i>Description</i>
Naive	1.0→0.0	1.e−8	
Umfpack	0.0→1.0	0.1	
KLU	0.0→1.0	0.1	
Y12m	boolean: none=0.0, full=1.0	full	
Lapack	1.0→0.0		
SuperLU	0.0→1.0	1.0	
TAUCS			
Watson			

matrix handling, `cc` or `dir`). However, this behavior can be influenced by the `threads` directive. The value `auto` is the default; the value `disable` reverts to the usual scalar behavior; otherwise `threads`, the desired number of threads, is read and used. If it is prefixed by the keyword `assembly`, the desired number of threads is used only in assembly; if it is prefixed by the keyword `solver`, the desired number of threads is used only during factorization/solution, if the linear solver supports it (currently, only `naive` and `superlu` solvers do). As an example, on a 4 CPU architecture, to use 2 threads for assembly and 4 threads for solution, use

```
threads: assembly, 2;
solver: superlu, cc, mt, 4;
```

4.1.2 Derivatives Data

Data related to the ‘derivatives’ phase. Right after the initial assembly and before the simulation starts, the so-called derivatives solution is performed. The system is solved with the kinematic unknowns constrained, in order to properly determine the dynamic unknowns, namely momenta and constraint reactions. For this purpose, the coefficient that relates the state perturbation to the derivative perturbation must be set to a value that is small enough to allow the determination of accurate derivatives with very small change in the states. This coefficient should be zero, but this leads to matrix singularity, so it must be chosen by the user, since it is highly problem dependent. A rule-of-thumb is: if the system has small stiffness and high inertia, the coefficient can be big, if the system has high stiffness and small inertia, the coefficient must be small.

The derivatives solution is always performed and cannot be disabled. If for any reason it does not converge, to practically disable it one can set a very large tolerance. Subsequent time steps may start with a less than ideal initialization, resulting in a rough transient.

Derivatives Tolerance

```
<card> ::= derivatives tolerance : <tolerance> ;
```

Derivatives Max Iterations

```
<card> ::= derivatives max iterations : <max_iterations> ;
```

Derivatives Coefficient

```
<card> ::= derivatives coefficient : <coefficient> ;
```

4.1.3 Dummy Steps Data

Data related to the ‘dummy steps’ phase. The dummy steps are intended as a sort of numerical computation of the second order derivatives of the constraint equations. The displacement constraint equations in an index three Differential-Algebraic Equations system (DAE) represent the second derivative effect of the constraints on the kinematic unknowns. Thus, to ensure the proper initial conditions, the constraint equations should be differentiated twice. To simplify the code, those equations have been differentiated once only, in the initial assembly, the second derivation being performed numerically by the dummy steps. During these steps the system converges to a solution whose error from the sought solution is bounded. For this reason, the dummy steps should be performed with a very short time step, to seek accuracy, and with a high numerical damping, to cancel as quickly as possible the high frequency numerical noise.

Note: no step change, modified Newton-Raphson and so on is allowed during the dummy steps.

Note: in most practical applications, dummy steps should not be used.

Dummy Steps Tolerance

```
<card> ::= dummy steps tolerance : <tolerance> ;
```

Dummy Steps Max Iterations

```
<card> ::= dummy steps max iterations : <max iterations> ;
```

Dummy Steps Number

```
<card> ::= dummy steps number : <number> ;
```

number of dummy steps.

Dummy Steps Ratio

```
<card> ::= dummy steps ratio : <ratio> ;
```

ratio of the time step to be used in the dummy steps to the regular time step.

Dummy Steps Method

```
<card> ::= dummy steps method : <method_data> ;
```

Same as for the normal simulation method.

4.1.4 Output Data

Model output, e.g. nodes and elements output, is handled by the data manager in the `control data` block; the solver only deals with simulation-related output, and takes care of some debug output.

Output

This command requests special output related to the solution phase.

```
<card> ::= output : <item> [ , ... ] ;
```

```
<item> ::=  
  { iterations  
    | residual  
    | solution  
    | jacobian matrix  
    | messages  
    | counter  
    | bailout  
    | none }
```

The keyword `counter` logs on the standard output a string that summarizes the statistics of the time step just completed. By default it is off; the same information, by default, is logged in the `.out` file. The keyword `iterations` logs on `stdout` a detailed output of the error at each iteration inside a time step. The keywords `residual`, `solution` and `jacobian matrix` log to `stdout` the residual, the solution and the Jacobian matrix; they are mainly intended for last resort debugging purposes, as the output can be extremely verbose. The item `messages` refers to all messaging on the `.out` file; by default, it is on. The

special item `bailout` instructs the solver to output the residual (as with `residual`) only in case of no convergence, before bailing out. The special item `none` clears the output flags; the items are additive, so, for instance, to clear out the default and add the iterations output, use:

```
output: none, iterations;
```

Output meter

This command conditions the output selected by the `output` command to a meter, expressed in terms of a drive caller. Only when the drive caller returns `true` (!0), the output occurs.

```
<card> ::= output meter : (DriveCaller) <when> ;
```

4.1.5 Real-Time Execution

Initially implemented by Michele Attolico

Extensively reworked by Pierangelo Masarati

This statement forces MBDyn to be scheduled in real-time, based on the capabilities of the underlying OS.

```
<card> ::= real time : [ { RTAI | POSIX } , ]
    <mode>
    [ , reserve stack , <stack_size> ]
    [ , allow nonroot ]
    [ , cpu map , <cpu-map> ]
    [ , output , { yes | no } ]
    [ , hard real time ]
    [ , real time log [ , file name , <command_name> ] ]

<mode> ::=
    { [ mode , period , ] time step , <time_step>
      | mode , semaphore [ , <semaphore_args> ]
      | mode , IO }
```

where:

- the optional `RTAI` or `POSIX` keywords indicate whether RTAI (Linux Real-Time Application Interface, <http://www.rtai.org/>) or POSIX should be used to schedule the execution of MBDyn in real-time (defaults to `RTAI`);
- when the keyword `mode` is set to `period`, the execution is scheduled periodically; when `mode` is set to `semaphore`, the execution waits for an external trigger to start a new step; when `mode` is set to `IO`, the execution assumes synchronization to be provided by waiting on blocking I/O streams (note: `semaphore` is not implemented yet, so `semaphore_args` is currently undefined; the default `mode` is `period`);
- the `time_step`, required when scheduling is `periodic`, is the sample rate in nanoseconds (ns); usually, it matches the time step of the simulation, but it can be different for other purposes;
- `allow nonroot` means that the program should be allowed to run as a regular (non-root) user while performing hard real-time operations with the underlying OS (by default, only the superuser is allowed to execute in real-time, since this requires the process to acquire the highest priority; only honored by `RTAI`);

- the keyword `reserve stack` instructs the program to statically reserve the desired stack size by means of the `mlockall(2)` system call; it should be used to ensure that no page swapping occurs during the real-time simulation; a minimal default value (1024 bytes) is set in any case;
- the keyword `cpu map` allows the program to force its execution on a specific subset of CPUs on multiprocessor hardware; the syntax of the `cpu_map` field is a byte-sized integer (between 0 and 255) whose active bits are the desired CPUs (up to 8);
- the keyword `output` determines whether output is entirely disabled or not; this option is only available to `POSIX` real-time, and should always be set to `no` in order to disable any output, so that I/O only occurs for the purpose of interprocess communication;
- the keyword `hard real time` instructs the program to run in hard real-time; the default is soft real-time (RTAI only; by default scheduling is in soft real-time);
- the keyword `real time log` enables logging by means of RTAI mailboxes; an optional `file name` argument allows to set the `command_name` of the logging process that is used. It should be a low priority soft real-time process that communicates with the simulator by means of the RTAI mailbox called `logmb`. The log process mailbox may receive a sequence of message consisting in two integers: the time step number and the amount of nanoseconds that time step overrun the scheduling period.
 Note: if no `command_name` is given, to guarantee the detection of the default log command the process `PATH` environment variable is augmented by prepending `.:BINPATH`; by default, `BINPATH` is the `bindir` directory of the build environment, so it can be set by using the `--bindir` configure switch.

The keywords must be given in the sequence reported above. Real-time simulation is mostly useless without interaction with external programs. The input is dealt with by the stream file drivers described in Sections 7.1.4 and 7.1.5. The output is dealt with by the stream output elements described in Sections 8.16.1, 8.16.2, and 8.16.3.

4.2 Other Problems

Other types of problems are either possible or being developed. Typically, static problems can be obtained in terms of a downgrading of the initial value problem, by eliminating time-dependent interactions and treating time as an ordering parameter.

A special problem that is currently under development is `inverse dynamics`.

Chapter 5

Control Data

This section is read by the manager of all the bulk simulation data, namely the nodes, the drivers and the elements. It is used to set some global parameters closely related to the behavior of these entities, to tailor the initial assembly of the joints in case of structural simulations, and to tell the data manager how many entities of every type it should expect from the following sections. Historically this is due to the fact that the data structure for nodes and elements is allocated at the beginning with fixed size. This is going to change, giving raise to a “free” and resizable structure. But this practice is to be considered reliable since it allows a sort of double-check on the entities that are inserted.

5.1 Assembly-Related Cards

The initial assembly related cards are:

5.1.1 Skip Initial Joint Assembly

```
<card> ::= skip initial joint assembly ;
```

This directive inhibits the execution of the initial joint assembly. Note that for a model to behave correctly, the initial joint assembly should always succeed. A correct model succeeds with 0 iterations, i.e. it intrinsically satisfies the constraints from the beginning. However, the initial joint assembly is more than a simple compliance test; it represents a static preprocessor for models. See the directives **use** in Section 5.1.2 and **initial stiffness** in Section 5.1.3 for more details on performing appropriate initial joint assembly.

5.1.2 Use

```
<card> ::= use : <item_list> , in assembly ;
```

```
<item_list> ::= <item> [ , ... ]
```

```
<item> ::=  
  { rigid bodies  
    | gravity  
    | forces  
    | beams  
    | aerodynamic elements  
    | loadable elements }
```

joints are used by default, and cannot be added to the list. **beams** are used by default, too, but can be added to the list essentially for backward compatibility.

5.1.3 Initial Stiffness

```
<card> ::= initial stiffness : <position_stiffness> [ , <velocity_stiffness> ] ;
```

This directive affects the stiffness of the dummy springs that constrain the position and the orientation (**position_stiffness**) and the linear and angular velocity (**velocity_stiffness**) of the structural nodes during the initial assembly; the default is 1.0 for both. Note that each node can use a specific value; see Section 6.1 for details.

Note that the same value is used for the position and the orientation, so this stiffness is dimensionally inconsistent; It should really be intended as a penalty coefficient. The same considerations apply to the penalty value for linear and angular velocities.

5.1.4 Omega Rotates

```
<card> ::= omega rotates : { yes | no } ;
```

Sets whether the imposed angular velocity should be considered attached to the node or fixed in the global system during the initial assembly.

5.1.5 Tolerance

```
<card> ::= tolerance : <tolerance> ;
```

The tolerance that applies to the initial joint assembly; this tolerance is used to test the norm 2 of the residual, because it is very important, for a correct start of the simulation, that the algebraic constraints be satisfied as much as possible. The alternate statement **initial tolerance** is tolerated for backwards compatibility.

5.1.6 Max Iterations

```
<card> ::= max iterations : <max_iterations> ;
```

The number of iterations that are allowed during the initial assembly. The alternate statement **max initial iterations** is tolerated for backwards compatibility.

5.2 General-Purpose Cards

5.2.1 Title

```
<card> ::= title : " <simulation_title> " ;
```

5.2.2 Print

```
<card> ::= print : <item> [ , ... ] ;
```

```
<item> ::=  
  { dof stats  
    | [ { dof | equation } ] description
```

```

| [ { element | node } ] connection
| all
| none }

```

- The keyword `dof stats` enables printing of all potential dof owner entities at initial assembly and at regular assembly, so that the index of all variables can be easily identified;
- the keyword `dof description` adds extra variable description;
- the keyword `equation description` adds extra equation description;
- the keyword `description` is a shortcut for both `dof description` and `equation description`;
- the keyword `element connection` prints information about element connectivity;
- the keyword `node connection` prints information about node connectivity;
- the keyword `connection` is a shortcut for both `node connection` and `element connection`;
- the keyword `all` enables all dof statistics printing;
- the keyword `none` disables all dof statistics printing (the default).

Note that, apart from `none`, the other values are additive, i.e. the statement

```

print: dof stats;
print: dof description;

```

is equivalent to

```

print: dof stats, dof description;

```

while the statement

```

print: none;

```

disables all.

5.2.3 Make Restart File

```

<card> ::= make restart file
        [ : { iterations , <iterations_between_restarts>
          | time , time_between_restarts } ] ;

```

The default (no arguments) is to make the restart file only at the end of the simulation.

Note: the `make restart file` statement is experimental and essentially abandoned.

5.2.4 Select Timeout

```

<card> ::= select timeout , { forever | <timeout> } ;

```

exit after `timeout` minutes when waiting for connections on `stream drives` or `stream output elements`. By default, no timeout is used, so `select(2)` waits forever.

5.2.5 Default Output

```
<card> ::= default output : <output_list> ;

<output_list> ::= { all | none | <output_item> } [ , ... ]

<output_item> ::=
  { reference frames
  | abstract nodes
  | electric nodes
  | hydraulic nodes
  | structural nodes [ | accelerations ]
  | aerodynamic elements
  | air properties
  | beams
  | electric elements
  | forces
  | genels
  | gravity
  | hydraulic elements
  | joints
  | rigid bodies
  | induced velocity elements }
```

Here the default output flag for a type of node or element can be set. It can be overridden for each entity either when it is created or later, for entity aggregates, in each entity module, by means of the **output** directive for **nodes** (see Section 6.6.1) and **elements** (see Section 8.17.4). Special values are:

- **all** enables output of all entities;
- **none** disables output of all entities;
- **reference frames** by default, reference frames are not output; when enabled, a special file **.rfm** is generated, which contains all the reference frames defined, using the syntax of the **.mov** file.
- **accelerations** enables output of linear and angular accelerations for dynamic structural nodes, which are disabled by default. Accelerations output can be enabled on a node by node basis; see **structural node** (see Section 6.1) for details.

Since values are additive, except for **none**, to select only specific entities use **none** first, followed by a list of the entities whose output should be activated.

Example.

```
begin: control data;
  # ...
  # disable all except structural nodes
  default output: none, structural nodes;
  # ...
end: control data;
```

5.2.6 Output Frequency

This statement is intended for producing partial output.

```
<card> ::= output frequency : <steps> ;
```

Despite the perhaps misleading name, this statement causes output to be produced every `steps` time steps, starting from the initial time. A more general functionality is now provided by the `output meter` statement (Section 5.2.7).

5.2.7 Output Meter

A drive that causes output to be generated when different from zero, while no output is created when equal to zero. It is useful to reduce the size of the output file during analysis phases that are not of interest.

```
<card> ::= output meter : (DriveCaller) <meter> ;
```

The functionality of the deprecated `output frequency` statement can be reproduced by using the `meter` drive caller as follows:

```
output meter: meter, 0., forever, steps, 10;
```

When integrating with variable time step, one may want to use the `closest next DriveCaller` (see Section 2.5.2); for example,

```
output meter: closest next, 2., forever, const, 0.01;
```

causes output to occur at times greater than or equal to multiples of 0.01 s, starting at 2 s.

5.2.8 Output Precision

Sets the desired output precision for those file types that honor it (currently, all the native output except the `.out` file). The default is 6 digits; since the output is in formatted plain text, the higher the precision, the larger the files and the slower the simulation.

```
<card> ::= output precision : <number_of_digits> ;
```

This will no longer be an issue when the binary output is implemented; in that case, the output format will likely be fixed (float), or an optional extended format (double) will be allowed.

5.2.9 Output Results

This deprecated statement was intended for producing output in formats compatible with other software. See Appendix C for a description of the types of output that MBDyn can provide. Most of them are produced in form of postprocessing, based on the default raw output.

Right now, the `output results` statement is only used to enable the experimental support for NetCDF output, which eventually will replace the current textual output:

```
<card> ::= output results : netcdf ;
```

5.2.10 Default Orientation

This statement is used to select the default format for orientation output. For historical reasons, MBDyn always used the 123 form of Euler angles (also known as Cardano angles). This statement allows to enable different formats:

```
<card> ::= default orientation : <orientation_type> ;

<orientation_type> ::=
  { euler123
    | euler313
    | euler321
    | orientation vector
    | orientation matrix }
```

where

- **euler123** is the historical representation by means of three angles that represent three consecutive rotations about axes 1, 2 and 3 respectively, always applied to the axis as it results from the previous rotation;
- **euler313** is similar to **euler123**, but the rotations are about axes 3, 1 and 3 in the given sequence. This is the usual form for Euler angles.
- **euler321** is similar to **euler123**, but the rotations are about axes 3, 2 and 1 in the given sequence.
- **orientation vector** is the vector whose direction indicates the axis of the rotation that produces the orientation, and whose modulus indicates the magnitude of that rotation;
- **orientation matrix** indicates the orientation matrix itself.

The default remains **euler123**.

Note: this change implies that by default the selected way will be used to represent orientations in input and output. This flag is not fully honored throughout the code, yet. Right now, only **structural nodes** and selected elements can output orientations as indicated by **default orientation**. However, there is no direct means to detect what format is used in the `.mov` file (while it is easy, for example, in the `.nc` file generated by NetCDF). As a consequence, it is the user's responsibility to keep track of what representation is being used and treat output accordingly.

5.2.11 Default Aerodynamic Output

This statement is used to select the default output of built-in aerodynamic elements.

```
<card> ::= default aerodynamic output :
  <custom_output_flag> [ , ... ] ;

<custom_output_flag> ::=
  { position
    | orientation [ , <orientation_description> ]
    | velocity
    | angular velocity
    | configuration # ::= position, orientation, velocity, angular velocity
    | force
```

```

| moment
| forces      # ::= force, moment
| all }      # equivalent to all the above

```

```

<orientation_description> ::= orientation description , <orientation_type>

```

The `orientation_type` is defined in Section 5.2.10.

Flags add up to form the default aerodynamic element output request. Flags may not be repeated.

Output occurs for each integration point. The kinematics refers to location, orientation, and linear and angular velocity of the reference point. The forces are actual forces and moments contributing to the equilibrium of a specific node. By default, only force and moment are output. The custom output is only available in NetCDF format; see Section A.1.4.

5.2.12 Default Beam Output

This statement is used to select the default output of beam elements.

```

<card> ::= default beam output : <custom_output_flag> [ , ... ] ;

```

```

<custom_output_flag> ::=
{ position
  | orientation [ , <orientation_description> ]
  | configuration      # ::= position, orientation
  | force
  | moment
  | forces             # ::= force, moment
  | linear strain
  | angular strain
  | strains           # ::= linear strain, angular strain
  | linear strain rate
  | angular strain rate
  | strain rates      # ::= linear strain rate, angular strain rate
  | all }            # equivalent to all the above

```

```

<orientation_description> ::= orientation description ,
  <orientation_type>

```

The `orientation_type` is defined in Section 5.2.10.

Flags add up to form the default beam output request. Flags may not be repeated. Output refers to the beam’s “evaluation points”. Strain rates are only available from viscoelastic beams; even if set, elastic beams will not output them.

By default, only forces are output, to preserve compatibility with the original output format. The custom output is only available in NetCDF format; see Section A.1.4.

5.2.13 Default Scale

```

<card> ::= default scale : <scale_list> ;

```

```

<scale_list> ::= <scale_pair> [ , <output_list> ]

```

```

<scale_pair> ::= { all | none | <scale_item> } , <scale_factor>

```

```

<scale_item> ::=
  { abstract nodes
    | electric nodes
    | hydraulic nodes
    | structural nodes
    | thermal nodes }

```

Define the residual scaling factor for all dof owners, or for specific types of dof owners. In principle, all dof owners should allow to define a scale factor, since they define a set of equations. In practice, only the above listed types support this option.

Residual scaling is only active when specifically requested by the related option for the `tolerance` keyword in the problem-specific section. See Section 4.1.1 for further details.

5.2.14 Model

```

<card> ::= model : <model_type> ;

<model_type> ::= static

```

This statement allows to set the model type to `static`, which means that all dynamic structural nodes will be treated as static, and inertia forces ignored. Gravity and centripetal acceleration will only be considered as forcing terms. See the `structural` node (Section 6.1) for details.

5.2.15 Rigid Body Kinematics

```

<card> ::= rigid body kinematics : <rbk_data> ;

<rbk_data> ::= { <const_rbk> | <drive_rbk> }

<const_rbk> ::= const
  [ , position , (Vec3) <abs_position> ]
  [ , orientation , (OrientationMatrix) <abs_orientation> ]
  [ , velocity , (Vec3) <abs_velocity> ]
  [ , angular velocity , (Vec3) <abs_angular_velocity> ]
  [ , acceleration , (Vec3) <abs_acceleration> ]
  [ , angular acceleration , (Vec3) <abs_angular_acceleration> ]

<drive_rbk> ::= drive
  [ , position , (TplDriveCaller<Vec3>) <abs_position> ]
  [ , orientation , (TplDriveCaller<Vec3>) <abs_orientation_vector> ]
  [ , velocity , (TplDriveCaller<Vec3>) <abs_velocity> ]
  [ , angular velocity , (TplDriveCaller<Vec3>) <abs_angular_velocity> ]
  [ , acceleration , (TplDriveCaller<Vec3>) <abs_acceleration> ]
  [ , angular acceleration , (TplDriveCaller<Vec3>) <abs_angular_acceleration> ]

```

In principle, the kinematic parameters should be consistent. However, in most cases this is not strictly required, nor desirable. In fact, if the model is made only of rigid bodies, algebraic constraints and deformable restraints, the case of a system rotating at constant angular velocity does not require `abs_angular_velocity` to be the derivative of `abs_orientation`, since the latter never appears in

the forces acting on the system. Similarly, the `abs_position` and `abs_velocity` do not appear as well, as soon as the latter is constant.

However, if other forces that depend on the absolute motion (position, orientation, and velocity) participate, this is no longer true. This is the case, for example, of aerodynamic forces, which depend on the velocity of the body with respect to the airstream, whose velocity is typically expressed in the global reference frame.

5.2.16 Loadable path

```
<card> ::= loadable path : [ { set | add } , ] " <path> " ;
```

This card allows to either set (optional keyword `set`) or augment (optional keyword `add`) the search path for run-time loadable modules using the path specified as the mandatory argument `path`.

Note: this command should be considered obsolete, but its replacement is not implemented yet. It impacts the loading of all run-time loadable modules, not only that of `user defined` or `loadable` elements (Section 8.15. See for example the `module load` statement in Section 2.3.6.

5.3 Model Counter Cards

The following counters can be defined:

5.3.1 Nodes

- `abstract nodes`
- `electric nodes`
- `hydraulic nodes`
- `parameter nodes`
- `structural nodes`
- `thermal nodes`

5.3.2 Drivers

- `file drivers`

5.3.3 Elements

- `aerodynamic elements`
- `aeromodals`
- `air properties`
- `automatic structural elements`
- `beams`
- `bulk elements`
- `electric bulk elements`

- electric elements
- external elements
- forces
- genels
- gravity
- hydraulic elements
- joints
- loadable elements
- output elements
- rigid bodies
- induced velocity elements

Chapter 6

Nodes

The `nodes` section is enclosed in the cards:

```
begin : nodes ;  
    # ...  
end : nodes ;
```

Every node card has the format:

```
<card> ::= <node_type> : <node_label>  
    <additional_args>  
    [ , scale , { default | <scale> } ]  
    [ , output , { yes | no | (bool)<output_flag> } ]  
    [ , <extra_arglist> ] ;
```

where `node_type` is one of the following:

- `abstract`
- `electric`
- `hydraulic`
- `parameter`
- `structural`
- `thermal`

The data manager reads the node type and the label and checks for duplication. If the node is not defined yet, the appropriate read function is called, which parses the rest of the card and constructs the node.

The optional `scale` keyword is only supported by

- `abstract`
- `electric`
- `hydraulic`
- `structural`
- `thermal`

nodes. See Section 5.2.13 for further details.

6.1 Structural Node

The 6 dof structural node can be `static`, `dynamic`, `modal` or `dummy`.

6.1.1 Static Node

The `static` keyword means no inertia is related to that node, so it must be appropriately constrained or attached to elastic elements. Static nodes are useful when there is no need to apply inertia to them, thus saving 6 degrees of freedom.

6.1.2 Dynamic Node

The `dynamic` keyword means inertia can be attached to the node, so it provides linear and angular momenta degrees of freedom, and automatically generates the so-called `automatic structural` elements.

6.1.3 Modal Node

The `modal` node is basically a regular `dynamic` node that must be used to describe the rigid reference motion of a `modal` joint. See Section 8.12.32 for further details.

6.1.4 Syntax

```
<additional_args> ::= , { static | dynamic | modal } ,
  (Vec3) <absolute_position> ,
  (OrientationMatrix) <absolute_orientation_matrix> ,
  [ orientation description , <orientation_description> , ]
  (Vec3) <absolute_velocity> ,
  (Vec3) <absolute_angular_velocity>
  [ , assembly
    , (Scalar) <position_initial_stiffness>
    , (Scalar) <velocity_initial_stiffness>
    , (YesNo) <omega_rotates?> ]

<orientation_description> ::=
  { euler123 | euler313 | euler321
    | orientation vector | orientation matrix }
```

The stiffness parameters and the `omega_rotates?` flag override the default values optionally set by the `initial stiffness` and `omega rotates` keywords in the `control data` block. They are optional, but they must be supplied all together if at least one is to be input.

The `omega_rotates?` parameter determines whether the initial angular velocity should follow or not the node as it is rotated by the initial assembly procedure. It can take values `yes` or `no`; a numerical value of 0 (no) or 1 (yes) is supported for backward compatibility, but its use is deprecated.

The `dynamic` and the `modal` node types allow the optional output keyword `accelerations` after the standard node output parameters:

```
<extra_arglist> ::= accelerations , { yes | no | <value> }
```

to enable/disable the output of the linear and angular accelerations of the node. Since they are computed as a postprocessing, they are not required by the regular analysis, so they should be enabled only when strictly required.

Also note that accelerations may be inaccurate, since they are reconstructed from the momentum and the momenta moment derivatives through the inertia properties associated to dynamic nodes.

Also **dummy** nodes allow the keyword **accelerations**, but honor it only if provided by the nodes they are related to.

Accelerations output can be controlled by means of the **default output** statement (see Section 5.2.5).

If the **static** model type is used in the control data block, all dynamic structural nodes are actually treated as static. This is a shortcut to ease running static analyses without the need to modify each node of a dynamic model.

6.1.5 Dummy Node

The **dummy** structural node has been added to ease the visualization of the kinematics of arbitrary points of the system during the simulation. It does not provide any degrees of freedom, and it must be attached to another node.

Two dummy structural node types are available.

```
<additional_args> ::= , dummy , <base_node_label> , <type> , <dummy_node_data>
```

Dummy nodes take the label **base_node_label** of the node they are attached to, followed by **type**, the type of dummy node, possibly followed by specific data, **dummy_node_data**. The following types are available:

- **offset:**

```
<type> ::= offset
```

```
<dummy_node_data> ::=
  (Vec3)          <relative_offset> ,
  (OrientationMatrix) <relative_orientation_matrix>
  [ , orientation description , <orientation_description> ]
```

It outputs the configuration of a point offset from the base node.

- **relative frame:**

```
<type> ::= relative frame
```

```
<dummy_node_data> ::= <reference_node_label>
  [ , position , (Vec3) <reference_offset> ]
  [ , orientation , (OrientationMatrix) <reference_orientation_matrix> ]
  [ orientation description , <orientation_description> , ]
  [ , pivot node , <pivot_node_label>
  [ , position , (Vec3) <pivot_offset> ]
  [ , orientation , (OrientationMatrix) <pivot_orientation_matrix> ] ]
```

It outputs the configuration of the base node in the frame defined by the node of label **reference_node_label**, optionally offset by **reference_offset** and with optional relative orientation **reference_orientation_matrix**.

If a **pivot_node_label** is given, the relative frame motion is transformed as if it were expressed in the reference frame of the pivot node, optionally offset by **pivot_offset** and with optional relative orientation **pivot_orientation_matrix**.

Example.

```
set: real Omega = 1.;
structural: 1, static, null, eye, null, 0.,0.,Omega;
structural: 1000, dummy, 1, offset, 1.,0.,0., eye;
structural: 1001, dummy, 1, relative frame, 1000;
structural: 2000, dynamic,
    0.,0.,1.,
    1, 1.,0.,0., 2, 0.,0.,1.,
    null,
    null,
    accelerations, yes;
```

Output. Structural nodes generate two kinds of output files:

- the .mov file;
- the .ine file.

The .mov Output File.

The first refers to the kinematics of the node; its extension is .mov, and for each time step it contains one row for each node whose output is required. The rows contain:

1	the label of the node
2–4	the three components of the position of the node
5–7	the three Euler angles that define the orientation of the node
8–10	the three components of the velocity of the node
11–13	the three components of the angular velocity of the node
14–16	the three components of the linear acceleration of the dynamic and modal nodes (optional)
17–19	the three components of the angular acceleration of the dynamic and modal nodes (optional)

All the quantities are expressed in the global frame, except for the **relative frame** type of **dummy** node, whose quantities are, by definition, in the relative frame.

There exist two variants of this output. The output of the orientation can be modified by requesting either the three components of the Euler vector, or the nine components of the orientation matrix.

When the Euler vector is requested, columns 5 to 7 contain its components. Beware that it is not possible to discriminate between the original output and this format without knowing what output type was requested.

When the orientation matrix is requested, columns 5 to 13 contains the elements of the matrix, written row-wise, that is: $r_{11}, r_{12}, r_{13}, r_{21}, \dots, r_{33}$. Note that in this case the total column count, if accelerations are not requested, corresponds to that of a structural node with optional accelerations, so it might be hard to discriminate as well.

This ambiguity is resolved when results are output in NetCDF format. See Section A.1.3 for details.

Note: actually, the angles denoted as “Euler angles” are the three angles that describe a rotation made of a sequence of three steps: first, a rotation about global axis 1, followed by a rotation about axis 2 of the frame resulting from the previous rotation, concluded by a rotation about axis 3 of the frame resulting from the two previous rotations. To consistently transform this set of parameters into some other representation, see the tools `eu2rot(1)`, `rot2eu(1)`, `rot2eup(1)`, `rot2phi(1)`. The functions that compute the relationship between an orientation matrix and the set of three angles and vice versa are `MatR2EulerAngles()` and `EulerAngles2MatR()`, in `matvec3.h`.

The .ine Output File.

The second output file refers only to dynamic nodes, and contains their inertia; its extension is .ine.

For each time step, it contains information about the inertia of all the nodes whose output is required. Notice that more than one inertia body can be attached to one node; the information in this file refers to the sum of all the inertia related to the node.

The rows contain:

1	the label of the node
2-4	item the three components of the momentum in the absolute reference frame
5-7	item the three components of the momenta moment in the absolute reference frame, with respect to the coordinates of the node, thus to a moving frame
8-10	the three components of the derivative of the momentum
11-13	the three components of the derivative of the momentum moment

Private Data. The following data is available:

all structural nodes:

1. "X[1]" position in global direction 1
2. "X[2]" position in global direction 2
3. "X[3]" position in global direction 3
4. "x[1]" position in direction 1, in the reference frame of the node
5. "x[2]" position in direction 2, in the reference frame of the node
6. "x[3]" position in direction 3, in the reference frame of the node
7. "Phi[1]" orientation vector in global direction 1
8. "Phi[2]" orientation vector in global direction 2
9. "Phi[3]" orientation vector in global direction 3
10. "XP[1]" velocity in global direction 1
11. "XP[2]" velocity in global direction 2
12. "XP[3]" velocity in global direction 3
13. "xP[1]" velocity in direction 1, in the reference frame of the node
14. "xP[2]" velocity in direction 2, in the reference frame of the node
15. "xP[3]" velocity in direction 3, in the reference frame of the node
16. "Omega[1]" angular velocity in global direction 1
17. "Omega[2]" angular velocity in global direction 2
18. "Omega[3]" angular velocity in global direction 3
19. "omega[1]" angular velocity in direction 1, in the reference frame of the node
20. "omega[2]" angular velocity in direction 2, in the reference frame of the node
21. "omega[3]" angular velocity in direction 3, in the reference frame of the node

22. **"E[1]"** Cardan angle 1 (about global direction 1)
 23. **"E[2]"** Cardan angle 2 (about local direction 2)
 24. **"E[3]"** Cardan angle 3 (about local direction 3)
 25. **"E313[1]"** Cardan angle 1 (about global direction 3)
 26. **"E313[2]"** Cardan angle 2 (about local direction 1)
 27. **"E313[3]"** Cardan angle 3 (about local direction 3)
 28. **"E321[1]"** Cardan angle 1 (about global direction 3)
 29. **"E321[2]"** Cardan angle 2 (about local direction 2)
 30. **"E321[3]"** Cardan angle 3 (about local direction 1)
 31. **"PE[0]"** Euler parameter 0
 32. **"PE[1]"** Euler parameter 1
 33. **"PE[2]"** Euler parameter 2
 34. **"PE[3]"** Euler parameter 3
- dynamic nodes only:
35. **"XPP[1]"** acceleration in global direction 1
 36. **"XPP[2]"** acceleration in global direction 2
 37. **"XPP[3]"** acceleration in global direction 3
 38. **"xPP[1]"** acceleration in direction 1, in the reference frame of the node
 39. **"xPP[2]"** acceleration in direction 2, in the reference frame of the node
 40. **"xPP[3]"** acceleration in direction 3, in the reference frame of the node
 41. **"OmegaP[1]"** angular acceleration in global direction 1
 42. **"OmegaP[2]"** angular acceleration in global direction 2
 43. **"OmegaP[3]"** angular acceleration in global direction 3
 44. **"omegaP[1]"** angular acceleration in direction 1, in the reference frame of the node
 45. **"omegaP[2]"** angular acceleration in direction 2, in the reference frame of the node
 46. **"omegaP[3]"** angular acceleration in direction 3, in the reference frame of the node

Note: Euler parameters actually do not take into account the whole orientation of a node, since they are post-processed from the orientation matrix. As a consequence, they only parametrize the minimum norm orientation that yields the current orientation matrix of the node. The same applies to the orientation vector φ .

Note: if accelerations are requested using the **string** form, their computation is enabled even if it was not explicitly enabled when the node was instantiated. However, if the **index** form is used, their computation must have already been explicitly enabled.

Note: dummy nodes based on dynamic nodes inherit the capability to provide access to linear and angular accelerations.

6.2 Electric Node

```
<additional_args> ::= [ , value , <initial_value>
    [ , derivative , <derivative_initial_value> ] ]
```

Note: the keywords `value` and `derivative` have been introduced recently; `value` is not mandatory, resulting in a warning, while `derivative` is required. The same applies to the `abstract node` and to the `hydraulic node`; the latter is an algebraic node, so only `value` is allowed.

Private Data. The following data is available:

1. "x" value
2. "xP" derivative value

6.3 Abstract Node

```
<additional_args> ::= , { algebraic | differential }
    [ , value , <initial_value>
    [ , derivative , <derivative_initial_value> ] ]
```

Note: abstract nodes are ancestors of all scalar node types. Many `genel` and `electric` elements can be connected to `abstract` nodes as well, since they directly use the ancestor class.

Output. The value of abstract nodes is output with file extension `.abs`; for each time step the output of the required nodes is written. The format of each row is

- the label of the node
- the value of the node
- the value of the node derivative, when `differential`

Private Data. The following data is available:

1. "x" value
2. "xP" derivative value

6.4 Hydraulic Node

The hydraulic node represents the pressure at a given location of a hydraulic circuit. It is derived from the generic scalar algebraic node; as a consequence, it can be connected to all elements that operate on generic scalar nodes.

The syntax is

```
<additional_args> ::= [ , value , <initial_value> ]
```

Private Data. The following data is available:

1. "x" value

6.5 Parameter Node

NOTE: `parameter` nodes are essentially obsolete; in most cases their purpose can be better fulfilled by `element` and `node` drive callers.

```
<additional_args> ::= { [ , value , <initial_value> ]
  | , element
  | , sample and hold , (NodeDof) <signal> , <sample_time>
  | , beam strain gage , <y> , <z> }
```

The parameter node is derived from the class scalar algebraic node, but it is used in a rather peculiar way: it doesn't own any degree of freedom, so it does not participate in the solution; it is rather used as a sort of placeholder for those elements that require to be connected to a scalar node that is not otherwise significant to the analysis. Thanks to the availability of the `parameter` node, these elements do not need to be reformulated with a grounded node, while the parameter node value can be changed during the solution by several means, listed in the following.

Element. When the argument list starts with the keyword `element`, the parameter node expects to be bound to an element, and to access bulk element data (see the `bind` statement, Section 8.17.1).

Sample and Hold. When the argument list starts with the keyword `sample and hold`, followed by a `NodeDof` specification and a sample time, the parameter node contains the value of the input signal, namely the value of the node, for the duration of the `sample_time`. This may be useful to preserve the value of some signal across time steps.

Beam Strain Gage. When the argument list starts with the keyword `beam strain gage`, followed by the coordinates of a point on the section of a beam, the `parameter` node expects to be bound to a `beam` element, and to access the measure of the axial strain at point `x`, `y` in the section plane as a combination of section strains and curvatures:

$$\varepsilon = \nu_x + \mathbf{z} \cdot \kappa_y - \mathbf{y} \cdot \kappa_z, \quad (6.1)$$

where

- ν_x is the axial strain of the beam;
- κ_y is the bending curvature of the beam about the y axis;
- κ_z is the bending curvature of the beam about the z axis.

The span-wise location of the point where the strain is evaluated is set in the `bind` statement (see Section 8.17.1).

Note: measuring strains by means of derivatives of interpolated positions and orientations may lead to inaccurate results; force summation should be used instead.

6.6 Miscellaneous

6.6.1 Output

There is an extra card, that is used to modify the output behavior of nodes:

```
<card> ::= output : <node_type> , <node_list> ;
```

```
<node_list> ::= { <node_label> [ , ... ]  
  | range , <node_start_label> , <node_end_label> }
```

`node_type` is a valid node type that can be read in the `nodes` block. In case the keyword `range` is used, all nodes of that type with label between `node_start_label` and `node_end_label` are set.

*Note: if a node should never (**no**) or always (**yes**) be output, its output flag should be set directly on the node card. The global behavior of all the nodes of a type can be set from the **control data** block by adding the node type to the item list in the **default output** card. Then, the specific output flag of sets of nodes can be altered by means of the **output** card in the **nodes** block. This allows high flexibility in the selection of the desired output. The same remarks apply to the output of the elements.*

If `node_type` is `structural`, the optional keyword `accelerations` can be used right after the `node_type` to enable the output of the accelerations.

Chapter 7

Drivers

The `drivers` section is enclosed by the cards:

```
begin : drivers ;  
  # ...  
end : drivers ;
```

Every `driver` card has the format:

```
<card> ::= <driver_type> : <arglist> ;
```

At present the only type `driver_type` of drivers supported is `file`.

7.1 File Drivers

The `file` drivers are defined by the statement

```
file : <file_arglist> ;
```

A comprehensive family of `file` drivers is available; their syntax is:

```
<file_arglist> ::= <label> , <normal_arglist>
```

The following `file` drivers are supported:

7.1.1 Fixed Step

```
<normal_arglist> ::= fixed step ,  
  { count | <steps_number> } ,  
  <columns_number> ,  
  initial time , { from file | <initial_time> } ,  
  time step , { from file | <time_step> } ,  
  [ interpolation , { linear | const } , ]  
  [ { pad zeroes , { yes | no }  
    | bailout , { none | upper | lower | any } } , ]  
  " <file_name> "
```

If the keyword `count` is used instead of the `steps_number`, the number of time records provided by the file is obtained by counting the valid records.

If the keyword `from file` is used instead of either `initial_time` or `time_step`, the value is read from the file. In this case, comment lines of the form

```
# initial time: <initial_time>
# time step: <time_step>
```

are expected; the corresponding values are parsed and checked. A valid value provided in the input file always overrides any corresponding value provided in the corresponding comment line.

The value at an arbitrary time is interpolated from the available data. If the requested value is out of time bounds, zero is returned, unless `pad zeroes` is set to `no`, which means that the first or the last set of values is respectively used. As an alternative, if `bailout` is set to `upper`, `lower` or `any`, the simulation is stopped as soon as the time goes out of the respective "time" bounds.

If `interpolation` is `linear` (the default), a linear interpolation is used. Otherwise, if `interpolation` is `const`, the value at the beginning of the time step that includes the current time is used.

The file format is

```
# an arbitrary number of comment lines starting with '#'
#
# comment lines may provide special parameters
# like "initial time", "time step" as indicated above; for example
#
# initial time: 0
# time step: 0.01
#
# channel #1 channel #2 ... channel #n
  1.      2.      ... 100.
...

```

Example.

```
begin: drivers;
  file: 100,
    fixed step,
    100,      # n. steps
    1,       # n. channels
    0,       # initial_time
    0.01,    # time_step
    "input.dat";

  file: 200,
    fixed step,
    count,   # n. steps
    1,      # n. channels
    from file, # initial_time
    from file, # time_step
    "input.dat";
end: drivers;
```

7.1.2 Variable Step

```
<normal_arglist> ::= variable step ,
    <channels_number> ,
    [ interpolation , { linear | const } , ]
    [ { pad zeroes , { yes | no }
      | bailout , { none | upper | lower | any } , } , ]
    " <file_name> "
```

The same considerations of the `fixed step` type apply.

The file format is

```
# an arbitrary number of lines starting with '#'
#
# time channel#1 channel#2 ... channel#n
0. 1. 2. ... 100.
0.01 1.2 1.8 ... 90.
...
```

The number of available records is computed as the number of non-comment lines (lines not starting with a hash mark '#'). The first column contains the time, while the remaining `channels_number` columns contain the values of the input channels at that time. Time values must grow monotonically.

7.1.3 Socket

```
<normal_arglist> ::= socket ,
    <columns_number> ,
    [ initial values , <value #1> , ... , ]
    { local , " <file_name> "
      | [ port , <port_number> , ] (AuthenticationMethod) <authentication> }
```

The driver binds to a socket listening on port `port_number` (defaults to 9011) or on the named pipe `file_name`; at the beginning of each time step, in case of connection, the driver expects some input data in text format, consisting in an authentication token (if required). The authentication token is usually in the form

```
user: <user_name> NEWLINE
password: <user_password> NEWLINE
```

where `NEWLINE` is a literal 'newline' (i.e. '\n'). White-spaces may be significant in `user_name`, and surely are in `user_password`.

The identification is followed by a label token, in the form

```
label: <label> NEWLINE
```

indicating the column being edited, followed by the desired changes; the connection is terminated by a single mark followed by a newline:

```
. NEWLINE
```

The permitted operations, at present, are:

```
value: <value> NEWLINE
```

sets the value the drive will assume from the current step on

```
inc: { yes | no } NEWLINE
```

tells whether to switch **on** or **off** the increment mode, resulting in subsequent value commands being “set” rather than “add”

```
imp: { yes | no } NEWLINE
```

tells whether to switch **on** or **off** the impulse mode; when **on**, subsequent value commands to be applied for one step only. At present, impulse mode supersedes any incremental mode, namely the value of the drive is reset to zero after one step. This behavior may change in the future.

7.1.4 RTAI Mailbox

This special drive is a variant of the Socket Stream (Section 7.1.5; under development yet) that reads the input data from a RTAI mailbox in non-blocking mode. It is intended as a means of communication between different processes running in real-time. The syntax is:

```
<normal_arglist> ::= RTAI input ,  
    stream drive name , " <stream_name> " ,  
    [ create , { yes | no } , ]  
    [ host , " <host_name> " , ]  
    [ { [ non ] blocking } , [ ... ] ]  
    <columns_number> ;
```

where

- **stream_name** is the name of the mailbox (a unique string no more than 6 characters long);
- the **create** keyword determines whether the mailbox must be created or looked-up as already existing on the system;
- **host_name** is the name or the IP of the remote host where the mailbox resides; note that if this field is given, **create** must be set to **no** and the mailbox must be already defined on the remote host;
- the number of channels **columns_number** determines how many columns can be accessed via the **file** drive caller mechanism of Section 2.5.13, as for all the other **file** drivers.

This part of the program is rapidly evolving, so do please not expect too much documentation and backward compatibility.

Note: at present, these elements require that the simulation be run in real-time mode (see Section 4.1.5); future development will allow to emulate the use of these elements also when the simulation is not run in real-time, e.g. for modeling or model debugging purposes.

7.1.5 Stream

```
<normal_arglist> ::= stream ,  
    name , " <stream_name> " ,  
    create , { yes | no } ,  
    [ { local , " <path> " ,  
        | [ port , <port_number> , ] [ host , " <host_name> " , ] } ]  
    [ { [ no ] signal  
        | [ non ] blocking } , [ ... ] , ]  
    [ input every , <steps> , ]
```

```

[ receive first , { yes | no } , ]
[ timeout , <timeout> , ]
[ echo , " <echo_file_name> "
  [ , precision , <precision> ]
  [ , shift , <shift> ] , ]
<columns_number>
[ , initial values , <value #1> , ... ]

```

The **stream** drive allows MBDyn to receive streamed inputs from remote processes both during regular simulations and during real-time simulations under RTAI. If the simulation is run in real-time, it uses RTAI mailboxes, otherwise regular UNIX sockets are used, either in the **local** or in the **internet** namespace.

This drive type is intended to allow the development of real-time models by running regular simulations for the purpose of debugging the model and the control process without the overhead and the potential problems of running in real-time. Then the same model can be run in real-time without changing the details of the communication with the controller process.

Non real-time simulation During non real-time simulations, streams operate in blocking mode. The meaning of the parameters is:

- **stream_name** indicates the name the stream will be known as; it is mostly useless, and must be no more than 6 characters long, since it is only allowed for compatibility with RTAI's mailboxes;
- the instruction **create** determines whether the socket will be created or looked for by MBDyn; if **create** is set to **no**, MBDyn will retry for 60 seconds and then give up;
- the keyword **local** indicates that a socket in the local namespace will be used; if **create** is set to **yes**, the socket is created, otherwise it must exist.
- either of the keywords **port** or **host** indicate that a socket in the internet namespace will be used; if **create** is set to **yes**, **host_name** indicates the host that is allowed to connect to the socket; it defaults to any host (0.0.0.0); if **create** is set to **no**, **host_name** indicates what host to connect to; the default is localhost (127.0.0.1); the default port is 9011;
- the keyword **no signal** disables raising a SIGPIPE in case the stream is read after it was closed by the peer;
- the keyword **input every** allows to read new driver values every **steps** time steps;
- the keyword **receive first** allows to enable/disable receiving data at the initial time (defaults to **yes**, thus data is expected at the initial time, during derivatives);
- the keyword **timeout** allows to set a timeout for each read operation; the timeout is expressed in seconds as a floating point number, with a theoretical micro-second resolution (the actual resolution is left to the OS);
- the keyword **echo** allows to echo the streamed data into file **echo_file_name**, using the optional **precision** and a **shift** for the time, to make sure step changes are captured correctly (a small negative value is recommended); the resulting file is suitable to re-execute the simulation using a **variable step** file driver, using a **const** value for the **interpolation** option;
- the keyword **initial values**, followed by **columns_number** real values, allows to set the initial values of each channel; this is useful for example if the initial value of a channel is needed during initial assembly, before the first set of values is read from the peer.

If no socket type is specified, i.e. none of the `local`, `port` and `host` keywords are given, a socket is opened by default in the internet namespace with the default IP and port; the `create` keyword is mandatory.

Real-time simulation During real-time simulations, streams wrap non-blocking RTAI mailboxes. The meaning of the parameters is:

- the parameter `stream_name` indicates the name the stream will be known as in RTAI's resource namespace; it must no more than 6 characters long, and actually represents the mailbox name;
- the instruction `create` determines whether the mailbox will be created or looked for by MBDyn;
- the keyword `local` is ignored;
- the keyword `host` indicates that a mailbox on a remote host will be used; it is useless when `create` is set to `yes`, because RTAI does not provide the possibility to create remote resources; if none is given, a local mailbox is assumed;
- the keyword `port` is ignored.

The parameter `columns_number` determines how many channels will be used. A channel is a double typed number; a `stream drive` can read an arbitrary number of simultaneous channels.

Chapter 8

Elements

The `elements` section is enclosed in the cards:

```
begin : elements ;  
  # ...  
end : elements ;
```

Every element card has the following format:

```
<card> ::= <elem_type> :  
  <arglist>  
  [ , output , { yes | no | (bool)<output_flag> } ]  
  [ , <extra_arglist> ] ;
```

where `elem_type` is one of the following:

- structural elements:
 - `automatic structural`
 - `beam`
 - `body`
 - `couple`
 - `gravity`
 - `joint`
 - `joint regularization`
 - `plate`
- aerodynamic elements:
 - `aerodynamic beam2`
 - `aerodynamic beam3`
 - `aerodynamic body`
 - `aeromodal`
 - `aircraft instruments`
 - `air properties`

- induced velocity
- electric elements:
 - electric
- hydraulic elements:
 - hydraulic
- thermal elements:
 - TODO ...
- output elements:
 - RTAI output
 - stream output
 - stream motion output
- generic elements:
 - bind
 - bulk
 - force
 - genel
 - loadable
 - user defined
- miscellaneous element cards

In case of elements that can be instantiated only once, like the `gravity` or the `air properties` elements, the `arglist` does not contain any label; otherwise, a label is expected first, to allow for checks on duplicated elements, namely:

```
<arglist> ::= <label> , <normal_arglist>
```

The data manager reads the element type and the label and checks for duplication. If the element is not defined yet, the proper read function is called, which parses the rest of the card and constructs the element. The elements are read as follows.

8.1 Aerodynamic Elements

8.1.1 Aerodynamic Body/Aerodynamic Beam2/3

These elements share the description of the aerodynamics. The former assumes the aerodynamic surface to be rigid, and takes its configuration from a single node, while the latter respectively relies on a two or three-node beam and use the same interpolation functions of the beam to compute the configuration at an arbitrary point.

The input format is:

```

<elem_type> ::= { aerodynamic body | aerodynamic beam2 | aerodynamic beam3 }

<normal_arglist> ::= <connectivity> ,
  (Shape<1D>)      <surface_chord> ,
  (Shape<1D>)      <surface_aerodynamic_center> ,
  (Shape<1D>)      <surface_b_c_point> ,
  (Shape<1D>)      <surface_twist> ,
  [ tip loss , (Shape<1D>) <tip_loss> , ]
  <integration_points>
  [ , control , (DriveCaller) <control_drive> ]
  [ , <airfoil_data> ]
  [ , unsteady , { bielawa } ]
  [ , jacobian , { yes | no | <bool> } ]
  [ , <custom_output> ]

<extra_arglist> ::= { std | gauss | node }

```

where

```

<connectivity> ::= { <body_conn> | <beam2_conn> | <beam3_conn> }

# <elem_type> ::= aerodynamic body
<body_conn> ::= <node_label>
  [ , [ user defined ] induced velocity , <induced_velocity_label>
    [ , passive ] , ]
  (Vec3)      <relative_surface_offset> ,
  (OrientationMatrix) <relative_surface_orientation> ,
  (Scalar)    <surface_span>

# <elem_type> ::= aerodynamic beam2
<beam2_conn> ::= <beam2_label>
  [ , [ user defined ] induced velocity , <induced_velocity_label>
    [ , passive ] , ]
  (Vec3)      <relative_surface_offset_1> ,
  (OrientationMatrix) <relative_surface_orientation_1> ,
  (Vec3)      <relative_surface_offset_2> ,
  (OrientationMatrix) <relative_surface_orientation_2> ,

# <elem_type> ::= aerodynamic beam3
<beam3_conn> ::= <beam3_label>
  [ , [ user defined ] induced velocity , <induced_velocity_label>
    [ , passive ] , ]
  (Vec3)      <relative_surface_offset_1> ,
  (OrientationMatrix) <relative_surface_orientation_1> ,
  (Vec3)      <relative_surface_offset_2> ,
  (OrientationMatrix) <relative_surface_orientation_2> ,
  (Vec3)      <relative_surface_offset_3> ,
  (OrientationMatrix) <relative_surface_orientation_3> ,

<airfoil_data> ::=
  { naca 0012

```

```
| rae 9671
| [ theodorsen , ] c81 , <c81_data> }
```

and

```
<c81_data> ::=
  { <c81_label>
    | multiple , <airfoil_number> ,
      <c81_label> , <end_point>
      [ , ... ]
    | interpolated , <airfoil_number> ,
      <c81_label> , <position>
      [ , ... ] }
```

The `custom_output` optional data consists in

```
<custom_output> ::= custom output ,
  <custom_output_flag> [ , ... ]
```

The values of `custom_output_flag` are defined in Section 5.2.11.

Flags add up to form the custom output request. Flags may not be repeated. By default, only forces are output. The custom output is only available in NetCDF format; see Section A.1.4.

The field `induced_velocity` indicates the label of the `induced velocity` element that this element is linked to. This means that the element can get information about the induced velocity and should supply information about the forces it generates. If the keyword `user defined induced velocity` is used, then `induced_velocity` refers to a `user defined` element. Note: if the keyword `induced velocity` is used and no induced velocity element is found with the label `induced velocity` a `user defined` element with that label is looked up. If the optional

An arbitrary relative orientation and offset is allowed for all elements with respect to the nodes they are linked to. This means that the aerodynamic beam offsets refer to the position of the beam's nodes, and have nothing to do with offsets related to the structural beam element.

The `Shape<1D>` entities are used to compute the physical chord, aerodynamic center, velocity measurement point (the point where the kinematic boundary conditions are evaluated) and twist as functions of the dimensionless abscissa along the span.

The optional `tip loss` keyword allows to define an additional shape, `tip_loss`, whose value is used to scale the value of the normal aerodynamic force. By default, the scale factor is 1. The `tip_loss` shape should have a value comprised between 0 and 1.

In any case, the user must be aware of the fact that Gauss integration will be used over the span of the element. This consists in evaluating forces at specific spanwise stations. This assumes that the function to be integrated over the spanwise domain is regular. Sharp variations, like tip loss concentrated in the outmost 2% of a rotor blade span, might be entirely missed when using too little spanwise elements with too little integration points.

The span of the `aerodynamic body` element is set by the user; the offset vector points to the center-span of the element. The span of the `aerodynamic beam2` and `aerodynamic beam3` elements is computed based on the metric of the beam, from the first to the last node.

The aerodynamic center and the velocity measurement points are measured relative to the centerline of the elements, that is the line in direction 3 of the local frame from the end of the offset vector. This line is assumed to be at the 25% of the airfoil chord when steady aerodynamic coefficients are used (`unsteady_flag` = 0). The direction 1 is assumed to be the “reference” line of the airfoil, from the trailing edge to the leading edge (points “forward”), while direction 2 is normal to the other two and goes from the lower to the upper side of the airfoil (points “up”). Figure 8.1 shows the arrangement of the airfoil geometry and properties.




Figure 8.1: Airfoil geometry

The `airfoil_data` defaults to a built-in NACA 0012 semi-analytical model (FIXME: the unsteady correction is buggy; use the `c81` mode instead).

The `multiple` mode of the `c81` data allows to specify more than one airfoil for an aerodynamic element; the transition between airfoils is sharp. The integer `airfoil_number` indicates how many airfoils are expected; the real `end_point` indicates where the influence zone for that airfoil ends, expressed in terms of a non-dimensional abscissa spanning $(-1, 1)$ along the reference line, roughly along axis 3 of the aerodynamic reference frame; `end_point` must not lie outside the element. So, for example, if airfoil NACA 0015 is used in the leftmost part of an element up to 1/4 span, NACA 0012 is used from 1/4 to 3/4 span, and NACA 0009 is used in the remaining rightmost 1/4, the syntax is:

```
set: integer naca0015 = 15;
set: integer naca0012 = 12;
set: integer naca0009 = 9;
c81 data: naca0015, "naca0015.c81";
c81 data: naca0012, "naca0012.c81";
c81 data: naca0009, "naca0009.c81";
# beginning of aerodynamic element definition...
multiple, 3,
    naca0015, -0.5,    # from -1.0 to -0.5
    naca0012, 0.5,    # from -0.5 to 0.5
    naca0009, 1.0,    # from 0.5 to 1.0
# ...rest of aerodynamic element definition
```

The `interpolated` mode of the `c81` data allows to specify a smooth transition between different airfoils inside an element. The interpolation occurs at the integration points where the aerodynamic data is required, and it is performed once for all at the beginning of the analysis. Since this operation is time consuming, and essentially unnecessary, the interpolated data can be generated once for all with the utility `util/c81merge` once the position of the integration point is known, and the `multiple` mode can be used to directly provide the interpolated data to the aerodynamic element.

The `theodorsen` aerodynamic data uses C81 data, and superimposes Wagner's approximation of the Theodorsen incompressible unsteady correction of 2D lift and moment coefficients. It is experimental.

The `extra_arglist` allows to define the style of the output. The `std` style (the default), the `gauss` and the `node` styles are illustrated in the output section.

Output

Aerodynamic elements, both bodies and beams, write their output with file extension `.aer`; for each time step the required elements are output. In any case the label of the element is output first. Three

different formats are available: **std** (the default), **gauss** and **node**.

std (or Coefficients at Gauss points): the output consists in a set of 8 numbers for each block, that describe data at each Gauss integration point; multiple blocks for a single element are written on the same line. The format is:

- the angle of attack at that station, in degrees (namely, the angle between the component of the airfoil velocity, evaluated at the velocity measurement point, that in the airfoil plane and a reference line on the airfoil)
- the local yaw angle, in degrees (namely, the angle whose tangent is the ratio between the axial and the inplane components of the airfoil velocity)
- the local Mach number
- the local lift coefficient
- the local drag coefficient
- the local aerodynamic moment coefficient
- a private number
- another private number

When **aerodynamic beam2** and **aerodynamic beam3** elements are considered, the output is repeated for each portion of the beam; so, for example, a two-node beam is split in two portions, so the output contains $2 \times$ **integration_points** data blocks, while a three-node beam is split in three portions, so the output contains $3 \times$ **integration_points** data blocks.

node: the format is:

- the label of the node
- the three components of the force applied to the node
- the three components of the couple applied to the node

When **aerodynamic beam2** and **aerodynamic beam3** elements are considered, the output is repeated on the same line for each node the element is connected to.

gauss (or Forces at Gauss points): the output consists in the forces and moments per unit length at each Gauss integration point; the format is:

- the direction of the wind velocity relative to the element frame
- the lift,
- the drag,
- and the aerodynamic moment per unit length

When **aerodynamic beam2** **aerodynamic beam3** elements are considered, the output is repeated on the same line for each portion of beam.

8.1.2 Aeromodal Element

Note: implemented by Giuseppe Quaranta; documented by Alessandro Scotti.

This element is used to model an aerodynamic modal element, i.e. an unsteady aerodynamic model that inherits the structural motion from a `modal` element. Its definition is very similar to that of the pure modal element, but it also includes some data representing unsteady aerodynamics in the time domain through the residualization matrices. This element is defined as follows:

```
<elem_type> ::= aeromodal

<normal_arglist> ::=
  <reference_modal_joint> ,
  (OrientationMatrix) <orientation> ,
  <reference_chord> ,
  <number_of_aerodynamic_states> ,
  [ rigid , ]
  " <modal_matrices_file> "
```

With this formulation, anytime an aeromodal element is defined, the user needs to declare the number of modal aerodynamic elements in use in the `control data` section. An `air properties` card definition is also required.

The keyword `rigid` indicates that the generalized aerodynamic forces provided by the model include global forces and moments associated to the rigid body motion of the underlying modal element (FIXME: untested).

There is also an optional gust model which is totally undocumented; for further information, please contact the Author(s).

The `modal_matrices_file` file includes the state space model in form of matrices A , B , C , D_0 , D_1 and D_2 , according to the representation

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{q}$$
$$\mathbf{f} = \mathbf{q} \left(\mathbf{C}\mathbf{x} + \mathbf{D}_0\mathbf{q} + \frac{2V_\infty}{c}\mathbf{D}_1\dot{\mathbf{q}} + \left(\frac{2V_\infty}{c}\right)^2 \mathbf{D}_2\ddot{\mathbf{q}} \right)$$

where \mathbf{q} are the modal variables that describe the structural motion, and \mathbf{f} are the unsteady aerodynamic forces that apply to the structural dynamics equations.

The file is formatted as follows:

```
*** MATRIX A
(<na> x <na> coefficients)
*** MATRIX B
(<na> x <ns> coefficients)
*** MATRIX C
(<ns> x <na> coefficients)
*** MATRIX D0
(<ns> x <ns> coefficients)
*** MATRIX D1
(<ns> x <ns> coefficients)
*** MATRIX D2
(<ns> x <ns> coefficients)
```

Example.

```
aeromodal: WING, WING_JOINT,  
  eye,  
  131.25, 10, "ha145b.fea";
```

The `aeromodal` element is declared with the label `WING`. This element is attached to a `modal` joint named `WING_JOINT`. The orientation of the aerodynamic reference with respect to the nodal reference is here expressed by the identity matrix (`eye`). The aerodynamic element chord is 131.25 inches. This quantity must be consistent with the system chosen to define the whole model (SI, for example; in this case, British Units). The next field, 10, indicates the number of states needed to use the aerodynamic model. `ha145b.fea` is the name of the file that contains the state space model matrices, obtained with an approximation chosen by the user. In this particular case, a 10 states Padé approximation has been chosen. This example is taken from the Bisplinghoff Ashley Halfman (BAH) Jet Transport Wing cantilevered wing with modal aerodynamic frequency response, computed by a double-lattice method at Mach 0.0. Data were extracted from the MSC-NASTRAN aeroelastic example file, named `ha145b`, while the aerodynamic state-space fitting has been computed using a Padé polynomial approximation (by Pasinetti & Mantegazza). All quantities are expressed in inches and pounds.

8.1.3 Aircraft Instruments

```
<elem_type> ::= aircraft instruments  
  
<normal_arglist> ::= <aircraft_node>  
  [ , orientation , (OrientationMatrix) <relative_orientation> ]
```

The `aircraft_node` represents the aircraft; it is assumed that the “nose” of the aircraft is toward the positive x direction of the node, and the “top” of the aircraft is toward the positive z direction of the node. An optional orientation can be added to change the orientation of the aircraft with respect to the node. This is useful, for example, with helicopters, where conventionally the positive direction of the x axis is nose to tail.

The available measures are accessed during the simulation by defining appropriate `parameter` nodes, and by binding the `aircraft instruments` element private data to the nodes by means of the `bind` mechanism, or directly by means of the `element` drive (see Section 2.5.11).

Private Data The following data are available:

- `"airspeed"`: the airspeed as seen by the reference point on the aircraft, i.e. the combination of the airstream speed and of the node speed
- `"groundspeed"`: the absolute value of the projection of the node speed in the xy plane.
- `"altitude"`: the z component of the node position
- `"attitude"`: $\tan^{-1}(r_{31}, r_{11})$ (FIXME: better $\arcsin(r_{31})$?)
- `"bank"`: $\tan^{-1}(r_{32}, r_{22})$ (FIXME: better $\arcsin(r_{32})$?)
- `"turn"`: (not available yet)
- `"slip"`: (not available yet)
- `"verticalspeed"`: the z component of the node velocity

- **"angleofattack"**: the angle between the local z and x components of the velocity (airstream plus node), expressed in the reference frame of the aircraft
- **"heading"**: the angle between the x axis of the aircraft and the "north" (the global x axis) about the global z axis; note: heading wraps about South (+180 deg from East, -180 deg from West).
- **"longitude"**: (FIXME: not implemented yet)
- **"latitude"**: (FIXME: not implemented yet)

Note: this element is eXperimental.

8.1.4 Air Properties

The properties of the airstream are made of the physical properties of the air plus the description of the airstream velocity direction and amplitude. The former can be expressed in different forms, while the latter are based on three-dimensional drive callers, `TplDriveCaller<Vec3>`.

```
<arglist> ::=
  { (DriveCaller) <air_density> , (Scalar) <sound_celerity>
    | std , { { SI | British }
      [ , temperature deviation , <delta_T> ]
      | <p0> , (DriveCaller) <rho0> ,
        <T0> , <dT/dz> , <R> , <g0> , <z1> , <z2> }
      [ , reference altitude , <z0> ] } ,
  (TplDriveCaller<Vec3>) <air_speed>
  [ , gust , <gust_model> [ , ... ] ]
```

The first form consists in the bare input of the air density, in form of a drive caller, and of the sound celerity, e.g.:

```
air properties: 1.225, 340.,
               1.,0.,0., 150.;
```

The second form uses standard air properties, both in the international system (SI) or in British units. A temperature deviation and an altitude offset can be defined, e.g.:

```
air properties: std, SI, temperature deviation, -55,
               reference altitude, 1000.,
               1.,0.,0., 150.;
```

where standard properties in SI are used, with a temperature deviation of -55 K and a reference altitude of 1000 m. The air properties are computed based on the z position of the point where the **air properties** are requested (plus the optional altitude offset). The last possibility lets the user input all the parameters required to compute the **air properties** based on the z position of the point where they are requested, namely the reference pressure `p0`, the reference density `rho0`, the reference temperature `T0`, the initial temperature gradient `dT/dz`, the gas constant `R`, the initial gravity acceleration `g0`, the bottom and top altitudes of the null temperature gradient region `z1` and `z2`; e.g., for SI units:

```
air properties: std,
               101325.,      /* Pa, 1 bar */
               1.2250,      /* kg/m^3 */
               288.16,     /* K, ISA: 15 C */
               -6.5e-3,    /* K/m */
```

```

287.,          /* J/kgK */
9.81,          /* m/s^2 */
11000.,       /* m */
25000.,       /* m */
temperature deviation, -55,
reference altitude, 1000.,
1.,0.,0., 150.;

```

The air properties are defined according to the formulas

$$\left. \begin{aligned}
 T &= T_0 + dT/dz \cdot z \\
 p &= p_0 \left(\frac{T}{T_0} \right)^{-\frac{g_0}{dT/dz \cdot R}} \\
 \rho &= \rho_0 \left(\frac{T}{T_0} \right)^{-\left(\frac{g_0}{dT/dz \cdot R} + 1\right)}
 \end{aligned} \right\} z < z_1 \quad (8.1a)$$

$$\left. \begin{aligned}
 T &= T_0 + dT/dz \cdot z_1 \\
 p &= \left(p_0 \left(\frac{T}{T_0} \right)^{-\frac{g_0}{dT/dz \cdot R}} \right) e^{-g_0 \frac{z-z_1}{dT/dz \cdot R}} \\
 \rho &= \rho_0 \left(\frac{T}{T_0} \right)^{-\left(\frac{g_0}{dT/dz \cdot R} + 1\right)} e^{-g_0 \frac{z-z_1}{dT/dz \cdot R}}
 \end{aligned} \right\} z_1 \leq z < z_2 \quad (8.1b)$$

The case of $z > z_2$ is not currently handled; the formulas for $z \geq z_1$ are actually used. The value of z is computed by adding the z component of the position where the air properties are computed to the reference altitude z_0 .

The asymptotic air properties are characterized by `air_speed`, the `TplDriveCaller<Vec3>` of the air speed, expressed in the global reference frame. The possibility to vary the density and the airspeed using a driver has little physical meaning, especially the former; is intended as a practical means to gradually introduce the desired air properties, and the related airloads, in the analysis.

Gust

If the optional `gust` keyword is used, a gust model can be added. Note that a very elementary gust model, represented by a uniform change in airstream speed and direction can be implemented by using a time-dependent airstream drive.

Gusts can also be appended later to the `air properties` element by using the statement

```
gust : <gust_model> ;
```

Front 1D Gust The syntax of the `front 1D` gust model is:

```

<gust_model> ::= front 1D ,
    (Vec3)      <front_direction> ,
    (Vec3)      <perturbation_direction> ,
    (Scalar)    <front_velocity> ,
    (DriveCaller) <front_profile>

```

This model consists in a uniform front, defined as

$$v(\mathbf{x}, t) = ng(\mathbf{f} \cdot \mathbf{x} + V_{\text{ref}} \cdot t)$$

where

- \mathbf{v} is the velocity perturbation;
- \mathbf{x} is the position of the point whose airstream velocity is being computed;
- t is the current time;
- \mathbf{n} is the unit vector `perturbation_direction` that defines the direction of the velocity perturbation;
- $g(\cdot)$ is the function `front_profile` that defines the gust profile;
- \mathbf{f} is the unit vector `front_direction` that defines the direction of propagation of the front;
- V_{ref} is the velocity `front_velocity` of propagation of the front in direction \mathbf{f} .

As an example, a transverse cosine-shaped gust, with a wavelength of 100 m and a peak velocity of 5 m/s moving downstream at the airstream speed, 100 m/s, in standard air, is presented:

```
set: real waveLength = 100.; # m
set: real V_inf = 100.;      # m/s
set: real V_g = 5.;         # m/s
air properties: std, SI,
  1.,0.,0., const, V_inf, # reference airstream along X
  gust, front 1D,
    1.,0.,0.,          # front moving along X
    0.,0.,1.,         # gust along Z
    V_inf,             # front moving at V_inf
    cosine, 0., pi/waveLength, V_g/2., one, 0.;
```

Scalar Function Wind Profile The syntax of the `scalar function` wind profile, implemented as a gust model within the `air properties`, is:

```
<gust_model> ::= scalar function ,
  reference position , (Vec3) <X0> ,
  reference orientation , (OrientationMatrix) <R0> ,
  (ScalarFunction) <sf>
```

It yields a uniform velocity profile along the x axis of the reference orientation as a function of the z axis component of the relative position; namely, given the relative position

$$z = \mathbf{e}_3 \cdot (\mathbf{x} - \mathbf{X0}), \quad (8.2)$$

the velocity is

$$\mathbf{v} = \mathbf{e}_1 \cdot \mathbf{sf}(z), \quad (8.3)$$

where \mathbf{e}_i is the i -th axis of the reference orientation `R0`.

Power Law Wind Profile The syntax of the `power law` wind profile, implemented as a gust model within the `air properties`, is:

```
<gust_model> ::= power law ,
  reference position , (Vec3) <X0> ,
  reference orientation , (OrientationMatrix) <R0> ,
  reference elevation , <z_ref> ,
  reference velocity , (DriveCaller) <v_ref> ,
  exponent , <exponent>
```

It yields a uniform velocity profile along the x axis of the reference orientation as a function of the z axis component of the relative position; namely, given the relative position

$$z = \mathbf{e}_3 \cdot (\mathbf{x} - \mathbf{X0}), \quad (8.4)$$

the velocity is

$$\mathbf{v} = \mathbf{e}_1 \cdot \mathbf{v_ref} \left(\frac{z}{z_ref} \right)^{\text{exponent}}, \quad (8.5)$$

where \mathbf{e}_i is the i -th axis of the reference orientation `R0`. Typical values of `exponent` are about 0.1.

Logarithmic Wind Profile The syntax of the `logarithmic` wind profile, implemented as a gust model within the `air properties`, is:

```
<gust_model> ::= logarithmic ,
  reference position , (Vec3) <X0> ,
  reference orientation , (OrientationMatrix) <R0> ,
  reference elevation , <z_ref> ,
  reference velocity , (DriveCaller) <v_ref> ,
  surface roughness length , <z_0>
```

It yields a uniform velocity profile along the x axis of the reference orientation as a function of the z axis component of the relative position; namely, given the relative position

$$z = \mathbf{e}_3 \cdot (\mathbf{x} - \mathbf{X0}), \quad (8.6)$$

the velocity is

$$\mathbf{v} = \mathbf{e}_1 \cdot \mathbf{v_ref} \cdot \frac{\log(z/z_0) - \psi_m}{\log(z_ref/z_0) - \psi_m}, \quad (8.7)$$

where \mathbf{e}_i is the i -th axis of the reference orientation `R0`.

The surface roughness length describes the typical roughness of the surrounding surface. It is a very small number in case of smooth surface (e.g. 0.01m for grass), or 1/20 to 1/30 of the typical obstacle's size (e.g. 1m for woods).

Output The output occurs in the `.air` file, which contains:

- a fake label, always set to 1
- the air density
- the sound celerity
- the three components of the reference air speed with respect to the inertial reference frame

Private Data The following data is available:

- "vxinf" the x component of the airstream speed (without any gust contribution)
- "vyinf" the y component of the airstream speed (without any gust contribution)
- "vzinf" the z component of the airstream speed (without any gust contribution)
- "vinf" the module of the airstream speed (without any gust contribution)

8.1.5 Generic Aerodynamic Force

This element computes generalized aerodynamic forces parametrized on dynamic pressure and angle of attack and sideslip angle.

```

<elem_type> ::= generic aerodynamic force

<normal_arglist> ::= <node_label> ,
  [ position , <relative_position> , ]
  [ orientation , <relative_orientation> , ]
  [ reference surface , <reference_surface> , ]
  [ reference length , <reference_length> , ]
  [ alpha first , { no | yes } , ]
  { <data_file_specification> | reference , <gaf_data_label> }

<data_file_specification> ::= file ,
  [ { angle units , { radians | degrees }
    | scale angles , <angle_scale_factor> } , ]
  [ scale lengths , <length_scale_factor> , ]
  " <data_file_name> "

```

This element computes aerodynamic forces and moments \mathbf{f}_a , \mathbf{m}_a proportional to the dynamic pressure $q = \rho \mathbf{v}^T \mathbf{v} / 2$ and to empirical coefficients tabulated as functions of the angle of attack α and the sideslip angle β . The angles are computed from the relative airstream velocity \mathbf{v} expressed in the reference frame of the body. The velocity is computed at the reference point, optionally offset from the node `node_label` by the offset `relative_position` (defaults to `null`), in a reference frame optionally oriented from that of the node `node_label` by the orientation matrix `relative_orientation` (defaults to `eye`), namely

$$\mathbf{o} = (\text{OrientationMatrix}) \text{relative_position} \quad (8.8a)$$

$$\mathbf{R}_h = (\text{Vec3}) \text{relative_orientation} \quad (8.8b)$$

$$\mathbf{v} = \mathbf{R}_h^T \mathbf{R}_n^T (\dot{\mathbf{x}}_n + \boldsymbol{\omega}_n \times \mathbf{R}_n \mathbf{o}), \quad (8.8c)$$

where \mathbf{x}_n and \mathbf{R}_n are the position and orientation of the node, and $\dot{\mathbf{x}}_n$ and $\boldsymbol{\omega}_n$ are its linear and angular velocity. By default the angles are computed from the components of \mathbf{v} according to the formulas

$$\alpha = \text{atan} \left(\frac{v_3}{\|v_1\|} \right) \quad (8.9a)$$

$$\beta = -\text{atan} \left(\frac{v_2}{\text{sign}(v_1) \sqrt{v_1^2 + v_2^2}} \right); \quad (8.9b)$$

as a consequence, $-\pi/2 \leq \alpha \leq \pi/2$ and $-\pi \leq \beta \leq \pi$. If the optional keyword **alpha first** is set to **yes**, the alternative formulas

$$\alpha = \text{atan} \left(\frac{v_3}{\text{sign}(v_1) \sqrt{v_1^2 + v_2^2}} \right) \quad (8.10a)$$

$$\beta = -\text{atan} \left(\frac{v_2}{\|v_1\|} \right). \quad (8.10b)$$

are used; as a consequence, $-\pi \leq \alpha \leq \pi$ and $-\pi/2 \leq \beta \leq \pi/2$. In both cases, they assume that the body reference system is loosely aligned as:

- axis 1 goes from tail to nose of the aircraft;
- axis 2 points “to the right”;
- axis 3 points towards the ground when the aircraft is in level flight.

Force and moment are consistently applied to the node as

$$\mathbf{f}_n = \mathbf{R}_n \mathbf{R}_h \mathbf{f}_a \quad (8.11a)$$

$$\mathbf{m}_n = \mathbf{R}_n \mathbf{R}_h \mathbf{m}_a + \mathbf{R}_n (\mathbf{o} \times (\mathbf{R}_h \mathbf{f}_a)) \quad (8.11b)$$

Note: this element does not contribute to the Jacobian matrix of the problem; as such, it may deteriorate the convergence properties of the solution when aerodynamic forces dominate the problem. Moreover, it is not suitable for the direct computation of the eigenvalues of a problem about a steady solution.

Note: this element, by itself, may not be adequate to model the rigid body aerodynamics of an aircraft, since it does not account for the dependence of force and moment on angle rates.

Output The following output is available:

1. column 1: element label
2. column 2: the angle of attack α
3. column 3: the sideslip angle β
4. columns 4–6: the components of force \mathbf{f}_a in local x , y and z directions
5. columns 7–9: the components of moment \mathbf{m}_a in local x , y and z directions, about the reference point (node plus offset)
6. columns 10–12: the components of force \mathbf{f}_n in global x , y and z directions
7. columns 13–15: the components of moment \mathbf{m}_n in global x , y and z directions, about the node

Private Data The **generic aerodynamic force** element outputs the following private data:

- **"Fx"**, **"Fy"**, **"Fz"**: force components in the global reference frame;
- **"Mx"**, **"My"**, **"Mz"**: moment components in the global reference frame, about the node;
- **"fx"**, **"fy"**, **"fz"**: force components in the local reference frame;
- **"mx"**, **"my"**, **"mz"**: moment components in the local reference frame, about the reference point;
- **"alpha"**: angle of attack (in radian)
- **"beta"**: sideslip angle (in radian)

The values computed during the last residual assembly are returned.

Generic Aerodynamic Element Data

Data is stored in ASCII format in a file.

An arbitrary number of comment lines is allowed at the beginning. Comment lines start with either a percent ‘%’ or a hash mark ‘#’ in the first column. Their content is discarded until the end of the line.

The first non-comment line must contain two integers separated by whitespace. The integers represent the expected number of angle of attack and sideslip angle values, N_α and N_β .

Another arbitrary number of comment lines is allowed.

A set of $N_\alpha \cdot N_\beta$ lines is expected. No comments or empty lines are allowed in between. Each line contains:

- column 1: the angle of attack, α
- column 2: the sideslip angle, β
- columns 3-5: the force coefficients $f_{x/q}$, $f_{y/q}$, $f_{z/q}$
- columns 6-8: the moment coefficients $m_{x/q}$, $m_{y/q}$, $m_{z/q}$

Notes:

1. lines are sorted as follows: all values of α are defined for each value of β ; the same values of α are expected for each value of β ;
2. the angle ranges are $-\pi/2 \leq \alpha \leq \pi/2$ and $-\pi \leq \beta \leq \pi$; or, if **alpha first** is **yes**, they are $-\pi \leq \alpha \leq \pi$ and $-\pi/2 \leq \beta \leq \pi/2$;
3. the angles are expected in radians; use the mutually exclusive optional keywords **angle units**, to specify either **radians** or **degrees**, or **scale angles**, to specify the **angle_scale_factor**. The **angle_scale_factor** is the factor that when multiplied by the angle transforms it into radians; for example, if angles are provided in degrees then **angle_scale_factor** = $\pi/180$;
4. $q = 1/2\rho V^2$ is the local reference dynamic pressure, where V is the norm of the velocity at the reference point;
5. the coefficients express forces and moments in the reference frame attached to the body;
6. the coefficients are either expected in dimensional or non-dimensional form. In the former case, the force coefficients represent areas, while the moment coefficients represent volumes, since they need to be multiplied by the dynamic pressure to become forces and moments. In the latter case, they are pure numbers; a **reference_surface** and **reference_length** must be defined in the configuration of the corresponding **generic aerodynamic force** element. When dimensional coefficients are specified, they can be rescaled by using the optional keyword **scale lengths** to specify the **length_scale_factor**.

Example. The content of the file `example.dat` is

```
# This is an example of data for the "generic aerodynamic force" element
# 5 values of angle of attack (alpha) and 4 values of sideslip angle (beta)
# are provided
5 4
# alpha beta fx/q fy/q fz/q mx/q my/q mz/q
-90 -180 0 0 0 0 0 0
-20 -180 0 0 0 0 0 0
```

```

  0 -180 0 0 0 0 0 0
 20 -180 0 0 0 0 0 0
 90 -180 0 0 0 0 0 0
-90 -20 0 0 0 0 0 0
-20 -20 0 0 0 0 0 0
  0 -20 0 0 0 0 0 0
 20 -20 0 0 0 0 0 0
 90 -20 0 0 0 0 0 0
-90  20 0 0 0 0 0 0
-20  20 0 0 0 0 0 0
  0  20 0 0 0 0 0 0
 20  20 0 0 0 0 0 0
 90  20 0 0 0 0 0 0
-90 180 0 0 0 0 0 0
-20 180 0 0 0 0 0 0
  0 180 0 0 0 0 0 0
 20 180 0 0 0 0 0 0
 90 180 0 0 0 0 0 0

```

The corresponding statement in the input file is

```

set: integer GAF_NODE = 10;
set: integer GAF_ELEM = 20;
generic aerodynamic force: GAF_ELEM, GAF_NODE,
    file, angle units, degrees, "example.dat";

```

8.1.6 Induced velocity

The **induced velocity** element is used to associate the aerodynamic elements that model the lifting surfaces of an aircraft, or the blades of a helicopter rotor, when some inflow related computations are required.

By means of different inflow models, and by means of the aerodynamic load contributions supplied by the aerodynamic elements, the **induced velocity** element is able to compute the induced velocity at an arbitrary point on the lifting surface or rotor disk. This velocity term in turn is used by the aerodynamic elements to determine a better estimate of the boundary conditions.

The syntax of the **induced velocity** elements is:

```

<elem_type> ::= induced velocity

<normal_arglist> ::= <induced_velocity_type> , <induced_velocity_data>

```

Rotor

Currently, induced velocity models are only implemented for helicopter rotors. Originally, this type of element was known as **rotor**, and the original syntax is preserved for backwards compatibility.

The syntax of the helicopter rotor **induced velocity** element is:

```

<induced_velocity_type> ::= rotor

<induced_velocity_data> ::= <craft_node> ,
    [ orientation , (OrientationMatrix) <rotor_orientation> , ]

```

```

<rotor_node> ,
induced velocity , <induced_velocity_model>

```

The optional `rotor_orientation` is required when axis 3 of the `craft_node` is not aligned with the rotor axis; axis 3 of the `rotor_node` must be aligned with the rotor axis.

There are five models of induced velocity. The first is no induced velocity; the syntax is:

```

<induced_velocity_model> ::= no

```

There is no argument list. This element does not compute any induced velocity, but still computes the rotor traction for output purposes, if output is required. The others have a fairly common syntax. The first three are `uniform`, `glauert` and `mangler` induced velocity models:

```

<induced_velocity_model> ::=
{ uniform | glauert [ , type , <glauert_type> ] | mangler } ,
<reference_omega> , <reference_radius>
[ , <option> [ , ... ] ]

```

```

<glauert_type> ::= { glauert | coleman | drees
| payne | white and blake | pitt and peters | howlett
| drees 2 }

```

```

<option> ::=
{ ground , <ground_node>
| delay , (DriveCaller) <memory_factor>
| max iterations , <max_iterations>
| tolerance , <tolerance>
| eta , <eta>
| correction , <hover_correction_factor>, <ff_correction_factor> }

```

- The `reference_omega` field is used to decide whether the induced velocity computation must be inhibited because the rotor speed is very low.
- The `reference_radius` field is used to make the rotor related parameters non-dimensional.
- The `ground` parameter is used to inform the rotor about the proximity to the ground; the z component of the distance between the rotor and the ground nodes, in the ground node reference frame (direction 3, positive), is used for an approximate correction of the axial inflow velocity [9].
- The `memory_factor`, the `hover_correction_factor` and the `ff_correction_factor` (forward flight) are used to correct the nominal induced velocity, according to the formula

$$U_{\text{effective}} = (1 - \text{memory_factor}) U_{\text{nominal}} + \text{memory_factor} U_{\text{previous}}$$

with

$$U_{\text{nominal}} = \frac{T}{2\rho AV_{\text{tip}} \sqrt{\frac{\lambda^2}{\text{hover_correction_factor}^4} + \frac{\mu^2}{\text{ff_correction_factor}^2}}}$$

The `delay` parameter is used to linearly combine the current reference induced velocity with the induced velocity at the previous step; no delay means there is no memory of the previous value.

The memory factor behaves like a discrete first-order low-pass filter. As a consequence, its behavior depends on the integration time step. The `memory_factor` parameter defaults to 0. The `hover_correction_factor` and `ff_correction_factor` parameters default to 1.

- The `max_iterations`, `tolerance` and `eta` parameters refer to the iteration cycle that is performed to compute the nominal induced velocity. After `max_iterations`, or when the absolute value of the difference between two iterations of the nominal induced velocity is less than `tolerance`, the cycle ends. Only a fraction `eta` of the difference between two iterations of the nominal induced velocity is actually used; `eta` defaults to 1. The default is to make only one iteration, which is backward-compatible with the original behavior.

The last induced velocity model uses a dynamic inflow model, based on [10], with 3 inflow states. The syntax is:

```
<induced_velocity_model> ::= dynamic inflow ,
    <reference_omega> ,
    <reference_radius>
    [ , <option> [ , ... ] ]

<option> ::=
    { ground , <<ground_node>
      | initial value , <const_vel> , <cosine_vel> , <sine_vel>
      | max iterations , <max_iterations>
      | tolerance , <tolerance>
      | eta , <eta>
      | correction , <hover_correction_factor> , <ff_correction_factor> }
```

Most of the parameters are the same as for the previous models. The optional `delay` parameter is no longer allowed. The three states, corresponding to uniform, fore-aft and lateral inflow, can be explicitly initialized by means of the optional `initial value` parameter.

Output The following output is available for all rotor elements:

1. column 1: element label
2. columns 2–4: rotor force in x , y and z directions (longitudinal, lateral and thrust components)
3. columns 5–7: rotor moment about x , y and z directions (roll, pitch and torque components)
4. column 8: mean inflow velocity, based on momentum theory
5. column 9: reference velocity at rotor center, sum of airstream and `craft_node` node velocity
6. column 10: rotor disk angle
7. column 11: advance parameter μ
8. column 12: inflow parameter λ
9. column 13: advance/inflow angle $\chi = \tan^{-1}(\mu/\lambda)$
10. column 14: reference azimuthal direction ψ_0 , related to rotor yaw angle
11. column 15: boolean flag indicating convergence in reference induced velocity computation internal iterations

12. column 16: number of iterations required for convergence

The **dynamic inflow** model adds the columns

13. column 17: constant inflow state

14. column 18: sine inflow state (lateral)

15. column 19: cosine inflow state (longitudinal)

Rotor force and moment (columns 2–4 and 5–7) are the aerodynamic force and moment exerted by the rotor aerodynamics on the **rotor_node**, projected in the reference frame of the **craft_node**, optionally modified by the **rotor_orientation** matrix. The conventional naming of longitudinal (or drag), lateral and thrust force, and roll, pitch and torque moment, refer to a rotorcraft whose x axis is the longitudinal (nose to tail) axis, whose y axis is the lateral (portside) axis, and whose z axis is the vertical (bottom to top) axis.

Private Data The following data is available:

1. **"Tx"** rotor force in x direction (longitudinal force)
2. **"Ty"** rotor force in y direction (lateral force)
3. **"Tz"** rotor force in z direction (thrust)
4. **"Mx"** rotor moment about x direction (roll moment)
5. **"My"** rotor moment about y direction (pitch moment)
6. **"Mz"** rotor moment about z direction (torque)

The rotor force and moment components are expressed in the same reference frame described in the Output Section above.

8.1.7 Rotor

Deprecated; see **induced velocity** (Section 8.1.6).

8.2 Automatic structural

The so called **automatic structural** element is automatically generated when a dynamic structural node is instantiated. As such, when defined in the **elements** block, the element already exists. The only reason to repeat its definition is to modify the values of the momentum and of the momenta moment, and to initialize their derivatives. The label must match that of the node it refers to.

```
<elem_type> ::= automatic structural  
  
<normal_arglist> ::=  
  (Vec3) <momentum> ,  
  (Vec3) <momenta_moment> ,  
  (Vec3) <momentum_derivative> ,  
  (Vec3) <momenta_moment_derivative>
```

All the provided values are recomputed during the initial derivatives phase, so they should be intended as initial values for the Newton iteration. In general, there is no need to provide this data; they can speed up initial convergence in case of systems that are not at rest in the initial configuration, with kinematic constraints that strongly affect the motion.

Private Data The following data is available:

1. "beta[1]" momentum in global direction 1
2. "beta[2]" momentum in global direction 2
3. "beta[3]" momentum in global direction 3
4. "gamma[1]" momenta moment in global direction 1
5. "gamma[2]" momenta moment in global direction 2
6. "gamma[3]" momenta moment in global direction 3
7. "betaP[1]" momentum derivative in global direction 1
8. "betaP[2]" momentum derivative in global direction 2
9. "betaP[3]" momentum derivative in global direction 3
10. "gammaP[1]" momenta moment derivative in global direction 1
11. "gammaP[2]" momenta moment derivative in global direction 2
12. "gammaP[3]" momenta moment derivative in global direction 3

8.3 Beam Elements

The family of finite volume beam elements implemented in MBDyn allows to model slender deformable structural components with a high level of flexibility.

The beam is defined by a reference line and by a manifold of orientations attached to the line. It is assumed that the direction 1 of the orientations lies along the reference line, but it is not strictly required to be tangent to it even in the reference configuration.

The beam element is defined by its nodes; currently, 2 and 3 node beam elements are implemented. Each node of the beam is related to a **structural node** by an offset and an optional relative orientation, to provide topological flexibility.

The beam element is modeled by means of an original Finite Volume approach [11], which computes the internal forces as functions of the straining of the reference line and orientation at selected points along the line itself, called *evaluation points*, which lie somewhere between two pairs of beam nodes.

At each evaluation point, a 6D constitutive law must be defined, which defines the relationship between the strains, the curvatures of the beam and their time derivatives and the internal forces and moments at the evaluation points.

The strains and curvatures and their time derivatives are obtained from the nodal positions and orientations by differentiating the interpolation functions.

The 6D constitutive laws are defined as

$$\begin{pmatrix} F_x \\ F_y \\ F_z \\ M_x \\ M_y \\ M_z \end{pmatrix} = \mathbf{f} \left(\begin{pmatrix} \varepsilon_x \\ \gamma_y \\ \gamma_z \\ \kappa_x \\ \kappa_y \\ \kappa_z \end{pmatrix}, \begin{pmatrix} \dot{\varepsilon}_x \\ \dot{\gamma}_y \\ \dot{\gamma}_z \\ \dot{\kappa}_x \\ \dot{\kappa}_y \\ \dot{\kappa}_z \end{pmatrix} \right)$$

where, if the convention of using x as beam axis is followed:

- F_x is the axial force component;
- F_y and F_z are the shear force components;
- M_x is the torsional moment component;
- M_y and M_z are the bending moment components;
- ε_x is the axial strain component;
- γ_y and γ_z are the shear strain components;
- κ_x is the torsional curvature component;
- κ_y and κ_z are the bending curvature component;
- f is an arbitrary function that defines the constitutive law.

8.3.1 Beam Section Constitutive Law

Typically, linear elastic or viscoelastic constitutive laws are used, although one may want to implement specific nonlinear elastic or elastic-plastic constitutive laws.

Beam Section Characterization

MBDyn allows the broadest generality in defining what a linear elastic constitutive law contains, since the entire 6×6 constitutive matrix can be input. This means that internal forces and moments can be arbitrarily related to generalized strains and curvatures. However, to make sense, a constitutive matrix at the section level, must satisfy some constraints, e.g. it is expected to be symmetric, although this is not strictly enforced by the code.

However, most of the info about the extra-diagonal terms of the stiffness matrix are not usually available. One easy way to work this around is to resort to any so-called composite beam section characterization analysis available in the literature.

For details, the reader is referred to [12] for a review of the topic, to [13] for an early work on the subject, and to [14] for a more recent review of the original formulation. The software that implements this type of analysis is called ANBA++. It is not free software, so far. Prospective users can contact the authors, through MBDyn developers.

Disclaimer

The following paragraphs are intended as a means to help users preparing data for MBDyn models in a consistent manner. By no means they indicate that the beam section stiffness properties must be provided in a specific reference frame. On the contrary, MBDyn allows as much generality as possible, and actually the variety of choices is redundant, since equivalent properties can be input in different ways.

This is intended to allow the code to suit the users' needs regardless of the original format of the input data. As such, all the transformations reported in the following are only intended as suggestions and should not be taken literally. For instance, rotations and offsets of reference points could be reversed, changing the values of the offsets, without affecting the final result.

The most important aspect of MBDyn notion of beam section properties is that the reference point and orientation, although arbitrary, must be unique, and the common notions of center of axial strain, shear center (and center of mass) have no special meaning.

Equivalent 6×6 Section of Isotropic Beam

When an isotropic beam section is considered, the 6×6 constitutive matrix, referred to an arbitrary point in the section, with an arbitrary orientation, can always be written in terms of elementary stiffness and geometrical properties. These are the properties that are usually available in tabular form either from simplified beam section analysis or by experiments. A sketch of a generic section is shown in Figure 8.3, where the arbitrary reference frame indicated by axes x , y and z originates from an arbitrary reference point on the section.

Isotropic uniform beam sections allow to group the internal forces and moments in two sets, together with their conjugated generalized strains: those related to shear stress and strain, and those related to axial stress and strain, as illustrated in Figure 8.2. There is no direct coupling between the two groups,

	ε_x	γ_y	γ_z	κ_x	κ_y	κ_z
F_x	A				A	A
F_y		S	S	S		
F_z		S	S	S		
M_x		S	S	S		
M_y	A				A	A
M_z	A				A	A

Figure 8.2: Constitutive coefficients grouping (S: shear, A: axial)

at the section level, so the corresponding coupling coefficients are always zero. This is no longer true when material anisotropy must be taken into account.

The 3×3 sub-blocks can be separately transformed in diagonal form by referring the corresponding properties to appropriate separate points in the beam section, and by applying an appropriate rotation about the axis of the beam.

Axial Stress and Strain Properties

This section considers the submatrix represented by the coefficients marked as A in Figure 8.2, under the assumption that it is symmetric.

First, the problem of obtaining axial stiffness properties referred to a generic point in a generic orientation is considered, when the properties referred to the axial strain center in the principal reference frame are known.

Then, the problem of extracting the location of the axial strain center and of the principal reference frame, and the principal bending stiffnesses from generic data is presented as well.

The two problems are complementary. Usually, the first one needs to be considered when engineering properties are available and the generic constitutive properties required by MBDyn need to be computed.

Diagonal to Generic Properties. First the transformation from diagonal to generic properties is considered. This transformation consists in rotating the section properties and then in referring them to a common reference point in the blade section.

The diagonal properties are described by the constitutive matrix

$$\begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^\dagger = \begin{bmatrix} EA & 0 & 0 \\ \text{sym.} & EJ_y & 0 \\ & & EJ_z \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}. \quad (8.12)$$

This constitutive matrix is expressed in a reference frame that is centered in the center of axial strain, indicated with the subscript as , and oriented according to the bending principal axes.

A rotation α about axis x is used to transform the properties into the common reference frame of the beam section. The internal forces and moments are thus transformed according to

$$\begin{aligned} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^* &= [R]_{\text{axial}} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^\dagger \\ &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{bmatrix} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^\dagger, \end{aligned} \quad (8.13)$$

while the strains and curvatures are transformed according to

$$\begin{aligned} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^* &= [R]_{\text{axial}}^{-T} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^\dagger \\ &= [R]_{\text{axial}} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^\dagger. \end{aligned} \quad (8.14)$$

As a consequence, the constitutive relationship becomes

$$\begin{aligned} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^* &= [R]_{\text{axial}} [A]^\dagger [R]_{\text{axial}}^T \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^* \\ &= \begin{bmatrix} EA & 0 & 0 \\ \text{sym.} & EJ_y \cos^2 \alpha + EJ_z \sin^2 \alpha & (EJ_y - EJ_z) \cos \alpha \sin \alpha \\ & (EJ_y - EJ_z) \cos \alpha \sin \alpha & EJ_z \cos^2 \alpha + EJ_y \sin^2 \alpha \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^*. \end{aligned} \quad (8.15)$$

An offset of the reference point results from the internal force and moment transformation

$$\begin{aligned} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix} &= [T]_{\text{axial}} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^* \\ &= \begin{bmatrix} 1 & 0 & 0 \\ z_{as} & 1 & 0 \\ -y_{as} & 0 & 1 \end{bmatrix} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^*. \end{aligned} \quad (8.16)$$

Similarly, the strains and curvatures are transformed according to the relationship

$$\begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix} = [T]_{\text{axial}}^{-T} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^*. \quad (8.17)$$

The constitutive relationship becomes

$$\begin{aligned}
\begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix} &= [T]_{\text{axial}} [A]^* [T]_{\text{axial}}^T \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix} \\
&= [T]_{\text{axial}} [R]_{\text{axial}} [A]^\dagger [R]_{\text{axial}}^T [T]_{\text{axial}}^T \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix} \\
&= \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ & A_{22} & A_{23} \\ \text{sym.} & & A_{33} \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}. \tag{8.18}
\end{aligned}$$

The values of the coefficients are

$$A_{11} = EA \tag{8.19a}$$

$$A_{12} = z_{as}EA \tag{8.19b}$$

$$A_{13} = -y_{as}EA \tag{8.19c}$$

$$A_{22} = EJ_y \cos^2 \alpha + EJ_z \sin^2 \alpha + z_{as}^2 EA \tag{8.19d}$$

$$A_{23} = (EJ_y - EJ_z) \cos \alpha \sin \alpha - y_{as}z_{as}EA \tag{8.19e}$$

$$A_{33} = EJ_z \cos^2 \alpha + EJ_y \sin^2 \alpha + y_{as}^2 EA \tag{8.19f}$$

Generic to Diagonal Properties. Consider now a generic axial portion of the constitutive properties, symmetric and usually positive-definite. The constitutive matrix can be transformed to diagonal form by moving the reference point by an offset in the plane of the section, and then by rotating the properties about axis x .

The offset is applied by the internal forces and moments transformation

$$\begin{aligned}
\begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^* &= [T]_{\text{axial}}^{-1} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix} \\
&= \begin{bmatrix} 1 & 0 & 0 \\ -z_{as} & 1 & 0 \\ y_{as} & 0 & 1 \end{bmatrix} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}. \tag{8.20}
\end{aligned}$$

The corresponding strains and curvatures transformation is

$$\begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^* = [T]_{\text{axial}}^T \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}. \tag{8.21}$$

The transformed constitutive relationship is

$$\begin{aligned}
\begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^* &= [T]_{\text{axial}}^{-1} [A] [T]_{\text{axial}}^{-T} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix} \\
&= \begin{bmatrix} A_{11} & A_{12} - z_{as}A_{11} & A_{13} + y_{as}A_{11} \\ & A_{22} - 2z_{as}A_{12} + z_{as}^2A_{11} & A_{23} - z_{as}A_{13} + y_{as}A_{12} - y_{as}z_{as}A_{11} \\ \text{sym.} & & A_{33} + 2y_{as}A_{13} + y_{as}^2A_{11} \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^*. \tag{8.22}
\end{aligned}$$

The location that decouples the axial force from the bending moments is

$$y_{as} = -\frac{A_{13}}{A_{11}} \quad (8.23a)$$

$$z_{as} = \frac{A_{12}}{A_{11}} \quad (8.23b)$$

When the location of Eq. (8.23) is considered, Eq. (8.22) becomes

$$\begin{aligned} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^* &= \begin{bmatrix} A_{11} & 0 & 0 \\ \text{sym.} & A_{22} - A_{12}^2/A_{11} & A_{23} - A_{12}A_{13}/A_{11} \\ & & A_{33} - A_{13}^2/A_{11} \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^* \\ &= \begin{bmatrix} A_{11}^* & 0 & 0 \\ \text{sym.} & A_{22}^* & A_{23}^* \\ & & A_{33}^* \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^*. \end{aligned} \quad (8.24)$$

When a rotation about axis x is considered, the internal forces and moments are transformed according to the relationship

$$\begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^\dagger = [R]_{\text{axial}}^T \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^*, \quad (8.25)$$

and the strains and curvatures are transformed according to

$$\begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^\dagger = [R]_{\text{axial}}^T \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^*. \quad (8.26)$$

The constitutive relationship becomes

$$\begin{aligned} \begin{Bmatrix} F_x \\ M_y \\ M_z \end{Bmatrix}^\dagger &= [R]_{\text{axial}}^T [A]^* [R]_{\text{axial}} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^\dagger \\ &= \begin{bmatrix} A_{11}^\dagger & 0 & 0 \\ \text{sym.} & A_{22}^\dagger & A_{23}^\dagger \\ & & A_{33}^\dagger \end{bmatrix} \begin{Bmatrix} \varepsilon_x \\ \kappa_y \\ \kappa_z \end{Bmatrix}^\dagger, \end{aligned} \quad (8.27)$$

with

$$A_{11}^\dagger = A_{11}^* \quad (8.28a)$$

$$A_{22}^\dagger = A_{22}^* \cos^2 \alpha + A_{33}^* \sin^2 \alpha + 2A_{23}^* \cos \alpha \sin \alpha \quad (8.28b)$$

$$A_{23}^\dagger = (A_{33}^* - A_{22}^*) \cos \alpha \sin \alpha + A_{23}^* (\cos^2 \alpha - \sin^2 \alpha) \quad (8.28c)$$

$$A_{33}^\dagger = A_{33}^* \cos^2 \alpha + A_{22}^* \sin^2 \alpha - 2A_{23}^* \cos \alpha \sin \alpha. \quad (8.28d)$$

The constitutive relationship is diagonal when $A_{23}^\dagger = 0$, namely

$$\begin{aligned} \alpha &= \frac{1}{2} \tan^{-1} \left(\frac{A_{22}^* - A_{33}^*}{2A_{23}^*} \right) \\ &= \frac{1}{2} \tan^{-1} \left(\frac{A_{11} (A_{22} - A_{33}) - A_{12}^2 + A_{13}^2}{A_{11}A_{23} - A_{12}A_{13}} \right). \end{aligned} \quad (8.29)$$

Shear Stress and Strain Properties

Consider now the submatrix represented by the coefficients marked as S in Figure 8.2, under the assumption that it is symmetric, as indicated in Equation (8.30):

$$\begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix} = \begin{bmatrix} S_{11} & S_{12} & S_{13} \\ & S_{22} & S_{23} \\ \text{sym.} & & S_{33} \end{bmatrix} \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix} \quad (8.30)$$

The orientation of the shear force components about the section axis can be selected in order to decouple them; by applying the transformation

$$\begin{aligned} \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix}^* &= [R_{\text{shear}}] \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix} \\ &= \begin{bmatrix} \cos \beta & -\sin \beta & 0 \\ \sin \beta & \cos \beta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix} \end{aligned} \quad (8.31)$$

The angle that decouples the shear forces is

$$\beta = \frac{1}{2} \tan^{-1} \left(\frac{2S_{12}}{S_{22} - S_{11}} \right)$$

representing a rotation about the axis x of the beam with respect to the origin of the initial reference frame as shown in Figure 8.3, and the resulting coefficients are

$$GA_y = S_{11} \cos^2 \beta + S_{22} \sin^2 \beta - 2S_{12} \sin \beta \cos \beta \quad (8.32)$$

$$GA_z = S_{11} \sin^2 \beta + S_{22} \cos^2 \beta + 2S_{12} \sin \beta \cos \beta \quad (8.33)$$

the shear block becomes

$$\begin{aligned} \begin{Bmatrix} F_y \\ F_x \\ M_x \end{Bmatrix}^* &= \begin{bmatrix} GA_y & 0 & S_{13} \cos \beta - S_{23} \sin \beta \\ & GA_z & S_{13} \sin \beta + S_{23} \cos \beta \\ \text{sym.} & & S_{33} \end{bmatrix} \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix}^* \\ &= \begin{bmatrix} GA_y & 0 & S_{13}^* \\ & GA_z & S_{23}^* \\ \text{sym.} & & S_{33} \end{bmatrix} \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix}^* \end{aligned}$$

The transformation of Equation (8.34) moves the point of application of the shear force of an arbitrary amount $\{y, z\}$ in the beam section, with respect to the reference frame rotated by β about the axis x of the beam, as indicated in Figure 8.3:

$$\begin{aligned} \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix}^\dagger &= [T_{\text{shear}}] \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix}^* \\ &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ z & -y & 1 \end{bmatrix} \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix}^* \end{aligned} \quad (8.34)$$

So the transformed shear block of the constitutive matrix becomes

$$\begin{aligned}
\begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix}^\dagger &= [T_{\text{shear}}] \begin{Bmatrix} F_y \\ F_z \\ M_x \end{Bmatrix}^* \\
&= [T_{\text{shear}}] [A] [T_{\text{shear}}]^T \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix}^\dagger \\
&= \begin{bmatrix} GA_y & 0 & S_{13}^* + zGA_y \\ & GA_z & S_{23}^* - yGA_z \\ \text{sym.} & & S_{33} - yS_{23}^* + zS_{13}^* + z(S_{13}^* + zGA_y) - y(S_{23}^* - yGA_z) \end{bmatrix} \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix}^\dagger
\end{aligned} \tag{8.35}$$

If the position of the point is selected in such a manner that the shear force and the torsional moment are decoupled, i.e., according to the definition of center of shear force (the point in a beam section where the application of a transverse force results in no twist)

$$\begin{aligned}
y &= \frac{S_{23}^*}{GA_z} \\
z &= -\frac{S_{13}^*}{GA_y}
\end{aligned}$$

the shear block becomes

$$\begin{aligned}
\begin{Bmatrix} F_y \\ F_x \\ M_x \end{Bmatrix}^\dagger &= \begin{bmatrix} GA_y & 0 & 0 \\ & GA_z & 0 \\ \text{sym.} & & S_{33} - S_{13}^{*2}/GA_y - S_{23}^{*2}/GA_z \end{bmatrix} \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix}^\dagger \\
&= \begin{bmatrix} GA_y & 0 & 0 \\ & GA_z & 0 \\ \text{sym.} & & GJ \end{bmatrix} \begin{Bmatrix} \gamma_y \\ \gamma_z \\ \kappa_x \end{Bmatrix}^\dagger
\end{aligned}$$

When the shear and torsional stiffnesses, and the position of the shear strain center and the orientation of the shear axes are available, the shear portion of the stiffness matrix can be computed by reversing the order of the transformations described in Equations (8.31–8.34), i.e.:

$$\begin{bmatrix} S_{11} & S_{12} & S_{13} \\ & S_{22} & S_{23} \\ \text{sym.} & & S_{33} \end{bmatrix} = [R_{\text{shear}}]^T [T_{\text{shear}}]^{-1} \begin{bmatrix} GA_y & 0 & 0 \\ 0 & GA_z & 0 \\ 0 & 0 & GJ \end{bmatrix} [T_{\text{shear}}]^{-T} [R_{\text{shear}}] \tag{8.36}$$

This expression implies that the stiffness properties are referred to an arbitrary point at $\{-y, -z\}$ from the shear center, in the shear reference frame, followed by a rotation into the section reference frame by an amount $-\beta$. The resulting coefficients are

$$\begin{aligned}
S_{11} &= GA_y \cos^2 \beta + GA_z \sin^2 \beta \\
S_{12} &= (GA_z - GA_y) \sin \beta \cos \beta \\
S_{13} &= yGA_z \sin \beta - zGA_y \cos \beta \\
S_{22} &= GA_z \cos^2 \beta + GA_y \sin^2 \beta \\
S_{23} &= yGA_z \cos \beta + zGA_y \sin \beta \\
S_{33} &= GJ + z^2GA_y + y^2GA_z
\end{aligned}$$

Note that the order of the rotation and reference point transportation is reversed with respect to the axial properties; this is mostly done for convenience in computing the coefficients, because the opposite would result in more complicated formulas; however, their development the other way 'round is straightforward.




Figure 8.3: Beam section

Generic Anisotropic Beam Section

When a generic anisotropic beam section is considered, the partitioning of axial and shear constitutive properties of Figure 8.2 is no longer possible. The axial and shear straining can be completely and arbitrarily coupled, resulting in a full constitutive matrix of the beam section,

$$\begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^\dagger = \begin{bmatrix} \mathbf{K}_{f\nu} & \mathbf{K}_{f\kappa} \\ \mathbf{K}_{m\nu} & \mathbf{K}_{m\kappa} \end{bmatrix}^\dagger \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^\dagger, \quad (8.37)$$

where \mathbf{f} and \mathbf{m} are the internal force and moment vectors, while ν and κ are the linear and angular strains. Usually, $\mathbf{K}_{m\nu} \equiv \mathbf{K}_{f\kappa}^T$, while $\mathbf{K}_{f\nu}$ and $\mathbf{K}_{m\kappa}$ are symmetric.

The reference point or the reference orientation of the constitutive matrix can be changed by a sequence of transformations consisting in a rotation and an offset.

The internal force and moment, after a rotation defined by the rotation matrix \mathbf{R} , become

$$\begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^* = \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^\dagger. \quad (8.38)$$

Similarly, the strains become

$$\begin{aligned} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^* &= \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix}^{-T} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^\dagger \\ &= \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^\dagger. \end{aligned} \quad (8.39)$$

As a consequence, the re-oriented constitutive relationship is

$$\begin{aligned} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^* &= \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix} \begin{bmatrix} \mathbf{K}_{f\nu} & \mathbf{K}_{f\kappa} \\ \mathbf{K}_{m\nu} & \mathbf{K}_{m\kappa} \end{bmatrix}^\dagger \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix}^T \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^* \\ &= \begin{bmatrix} \mathbf{R}\mathbf{K}_{f\nu}^\dagger\mathbf{R}^T & \mathbf{R}\mathbf{K}_{f\kappa}^\dagger\mathbf{R}^T \\ \mathbf{R}\mathbf{K}_{m\nu}^\dagger\mathbf{R}^T & \mathbf{R}\mathbf{K}_{m\kappa}^\dagger\mathbf{R}^T \end{bmatrix} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^* \\ &= \begin{bmatrix} \mathbf{K}_{f\nu} & \mathbf{K}_{f\kappa} \\ \mathbf{K}_{m\nu} & \mathbf{K}_{m\kappa} \end{bmatrix}^* \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^*. \end{aligned} \quad (8.40)$$

The internal moment, after considering an offset $\mathbf{o} = [0, y, z]^T$ of the reference point of the constitutive properties, becomes $\mathbf{m} = \mathbf{m}^* + \mathbf{o} \times \mathbf{f}^*$. The internal force and moment then becomes

$$\begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{o} \times & \mathbf{I} \end{bmatrix} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^*. \quad (8.41)$$

Similarly, the linear and angular strain vectors become

$$\begin{Bmatrix} \nu \\ \kappa \end{Bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{o} \times & \mathbf{I} \end{bmatrix}^{-T} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}^*. \quad (8.42)$$

As a consequence, the offset constitutive relationship becomes

$$\begin{aligned} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix} &= \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{o} \times & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{K}_{f\nu} & \mathbf{K}_{f\kappa} \\ \mathbf{K}_{m\nu} & \mathbf{K}_{m\kappa} \end{bmatrix}^* \begin{bmatrix} \mathbf{I} & \mathbf{o} \times^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix} \\ &= \begin{bmatrix} \mathbf{K}_{f\nu}^* & \mathbf{K}_{f\kappa}^* \\ \mathbf{o} \times \mathbf{K}_{f\nu}^* + \mathbf{K}_{m\nu}^* & \mathbf{o} \times \mathbf{K}_{f\kappa}^* + \mathbf{K}_{m\kappa}^* \end{bmatrix} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix} \\ &= \begin{bmatrix} \mathbf{K}_{f\nu} & \mathbf{K}_{f\kappa} \\ \mathbf{K}_{m\nu} & \mathbf{K}_{m\kappa} \end{bmatrix} \begin{Bmatrix} \nu \\ \kappa \end{Bmatrix}. \end{aligned} \quad (8.43)$$

It might be tempting to find what offset and rotation allows to decouple the force and moment constitutive properties. This can be sought by defining a generic transformation

$$\begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{T} & \mathbf{I} \end{bmatrix} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^\dagger, \quad (8.44)$$

such that

$$\begin{aligned} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix} &= \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{T} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{K}_{f\nu} & \mathbf{K}_{f\kappa} \\ \mathbf{K}_{m\nu} & \mathbf{K}_{m\kappa} \end{bmatrix}^\dagger \begin{bmatrix} \mathbf{I} & \mathbf{T}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{Bmatrix} \boldsymbol{\nu} \\ \boldsymbol{\kappa} \end{Bmatrix} \\ &= \begin{bmatrix} \mathbf{K}_{f\nu}^\dagger & \mathbf{K}_{f\nu}^\dagger \mathbf{T}^T + \mathbf{K}_{f\kappa}^\dagger \\ \mathbf{T} \mathbf{K}_{f\nu}^\dagger + \mathbf{K}_{m\nu}^\dagger & \mathbf{T} \mathbf{K}_{f\nu}^\dagger \mathbf{T}^T + \mathbf{K}_{m\nu}^\dagger \mathbf{T}^T + \mathbf{T} \mathbf{K}_{f\kappa}^\dagger + \mathbf{K}_{m\kappa}^\dagger \end{bmatrix} \begin{Bmatrix} \boldsymbol{\nu} \\ \boldsymbol{\kappa} \end{Bmatrix}. \end{aligned} \quad (8.45)$$

The expected decoupling results from

$$\mathbf{T} = -\mathbf{K}_{m\nu}^\dagger \left(\mathbf{K}_{f\nu}^\dagger \right)^{-1}. \quad (8.46)$$

However, the resulting transformation \mathbf{T} is not guaranteed to have the skew-symmetric structure of $\boldsymbol{\omega} \times$, thus the decoupling may not be reducible to an offset.

The reverse transformation is relatively straightforward, after noticing that, according to Eq. (8.41),

$$\begin{aligned} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}^* &= \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \boldsymbol{\omega} \times & \mathbf{I} \end{bmatrix}^{-1} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix} \\ &= \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\boldsymbol{\omega} \times & \mathbf{I} \end{bmatrix} \begin{Bmatrix} \mathbf{f} \\ \mathbf{m} \end{Bmatrix}. \end{aligned} \quad (8.47)$$

So, as expected, it is sufficient to revert the sign of the offset to revert the transformation. The signs of the constitutive relationship change accordingly.

Locking Correction for Two-Node Beam

The three-node finite volume element has been implemented first, and uses conventional polynomial parabolic interpolation of the nodal displacements and orientations; the two-node finite volume element has been introduced later. This latter element presents some shear-locking, which, for linear elastic constitutive laws, may be overcome by correcting the section stiffness matrix in a relatively straightforward form:

$$\hat{\mathbf{K}} = \left(\mathbf{F} + \frac{L^2}{12} \mathbf{T} \mathbf{F} \mathbf{T}^T \right)^{-1}, \quad (8.48)$$

where $\mathbf{F} = \mathbf{K}^{-1}$ is the compliance matrix of the section, L is the length of the beam, i.e. the distance between the two reference points obtained by adding the optional offset to the nodes, and

$$\mathbf{T} = \begin{bmatrix} \mathbf{0} & \mathbf{e}_x \times \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$$

is the ‘‘arm’’ matrix that appears in the differential equilibrium equation

$$\boldsymbol{\vartheta}_{/x} - \mathbf{T}^T \boldsymbol{\vartheta} + \boldsymbol{\phi} = 0,$$

where $\boldsymbol{\vartheta} = [\mathbf{f}^T \mathbf{m}^T]^T$ are the internal forces and moments, while $\boldsymbol{\phi}$ are the external forces and moments per unit span.




Figure 8.4: Geometry of the three-node beam element.

There are no provisions to automatically apply the correction when defining the constitutive law of the section. The two-node beam has been reimplemented using a helicoidal interpolation of the nodal positions and orientations, to improve its capability to undergo large displacements and relative rotations. It is activated by using the keyword `hbeam2` instead of `beam2`. However, to reduce the shear-locking effect, the stiffness properties still need to be manually corrected according to Equation (8.48). The `hbeam2` element is *experimental*, and should be used only for development purposes.

8.3.2 Three-node beam element

The three-node beam element is described in detail in [11]. Each of the three points the beam element connects is referred to a structural node but can have an arbitrary offset to allow high generality in locating the structural reference line of the beam. Figure 8.4 illustrates the geometry of the three-node beam element.

The finite volume formulation presented in [11] is used. As a consequence, the internal forces and moments are evaluated at two points that are at about midpoint between nodes 1 and 2, and nodes 2 and 3 (at $\xi = -1/\sqrt{3} \cong -0.57735$ and $\xi = 1/\sqrt{3} \cong 0.57735$ of a non-dimensional abscissa ξ running from $\xi = -1$ at node 1 to $\xi = 1$ at node 3).

So the constitutive properties must be supplied in these points, as well as the orientation matrices \mathbf{R}_I and \mathbf{R}_{II} , that express the orientation of the reference system the constitutive properties are expressed in with respect to the global frame (the axial force is conventionally defined in direction 1). Any of the supported 6D constitutive laws can be supplied to define the constitutive properties of each beam section.

The traditional input format is

```
<elem_type> ::= beam3

<normal_arglist> ::=
  <node_1_label> , (Vec3) <relative_offset_1> ,
  <node_2_label> , (Vec3) <relative_offset_2> ,
  <node_3_label> , (Vec3) <relative_offset_3> ,
  (OrientationMatrix) <orientation_matrix_section_I> ,
```

```

(ConstitutiveLaw<6D>) <constitutive_law_section_I> ,
{ same | (OrientationMatrix) <orientation_matrix_section_II> } ,
{ same | (ConstitutiveLaw<6D>) <constitutive_law_section_II> }
[ , <custom_output> ]

```

The `relative_offset_<i>`, with $i=1,2,3$, are the vectors \mathbf{o}_i , with $i=1,2,3$, of Figure 8.4.

The orientation matrices `orientation_matrix_section_<j>`, with $j=I,II$, are the section orientation matrices \mathbf{R}_j , with $j=I,II$, of Figure 8.4.

The first keyword `same`, alternative to the `orientation_matrix_section_II`, means that the same orientation defined for the first point will be used for the second point.

The second keyword `same`, alternative to the `constitutive_law_section_II`, means that the same constitutive law defined for the first point will be used for the second point.

If any of the constitutive laws is either viscous or viscoelastic, the viscoelastic variant of the beam element is used. Otherwise, the elastic variant is used.

A more complete input format is

```

<elem_type> ::= beam3

<normal_arglist> ::=
  <node_1_label> ,
  [ position , ] (Vec3) <relative_offset_1> ,
  [ orientation , (OrientationMatrix) <relative_orientation_1> , ]
  <node_2_label> ,
  [ position , ] (Vec3) <relative_offset_2> ,
  [ orientation , (OrientationMatrix) <relative_orientation_2> , ]
  <node_3> ,
  [ position , ] (Vec3) <relative_offset_3> ,
  [ orientation , (OrientationMatrix) <relative_orientation_3> , ]
  { (OrientationMatrix) <orientation_matrix_section_I>
    | from nodes } ,
  (ConstitutiveLaw<6D>) <constitutive_law_section_I> ,
  { same
    | (OrientationMatrix) <orientation_matrix_section_II>
    | from nodes } ,
  { same
    | (ConstitutiveLaw<6D>) <constitutive_law_section_II> }
  [ , <custom_output> ]

```

This format is a superset of the traditional one, which is extended by adding the possibility to set relative node orientations that can be subsequently used to interpolate the orientation matrices at the evaluation points, by providing the keyword `from nodes` instead of the matrix. If the keyword `same` is used for the second evaluation point, the same method is used to compute the orientation matrix.

The `custom_output` optional data consists in

```

<custom_output> ::= custom output , <custom_output_flag> [ , ... ]

```

The values of `custom_output_flag` are defined in Section 5.2.12.

Flags add up to form the custom output request. Flags may not be repeated. Strain rates are only available from viscoelastic beams. By default, only forces are output, to preserve compatibility with the original output format. The custom output is only available in NetCDF format; see Section A.1.4.

As an example, a simple beam element, with diagonal section stiffness matrix is presented:

```

set: integer beam_label = 1000;
set: integer beam_node1 = 2001;
set: integer beam_node2 = 2002;
set: integer beam_node3 = 2003;
set: real EA = 1e6; # N
set: real GAy = .6e6; # N
set: real GAz = .6e6; # N
set: real GJ = 1.e3; # Nm^2
set: real EJy = 2.e3; # Nm^2
set: real EJz = 1.e4; # Nm^2
beam3: beam_label,
      beam_node1, reference, node, null,
      beam_node2, reference, node, null,
      beam_node3, reference, node, null,
      eye,
      linear elastic generic, diag,
          EA, GAy, GAz, GJ, EJy, EJz,
      same,
      same;

```

A not-so-simple beam section, where the center of axial strain and the shear center are not coincident, is illustrated below. The node offset is used to align the reference line with the shear center, and the axial strain center offset is used in the constitutive matrix:

```

set: integer beam_label = 1000;
set: integer beam_node1 = 2001;
set: integer beam_node2 = 2002;
set: integer beam_node3 = 2003;
set: real EA = 1e6; # N
set: real GAy = .6e6; # N
set: real GAz = .6e6; # N
set: real GJ = 1.e3; # Nm^2
set: real EJy = 2.e3; # Nm^2
set: real EJz = 1.e4; # Nm^2
set: real yas = 2.e-2; # m
set: real zas = 1.e-2; # m
set: real ysc = 4.e-2; # m
set: real zsc = 2.e-2; # m
set: real y = yas-ysc; # compute the axial strain center
set: real z = zas-zsc; # wrt/ the shear center
beam3: beam_label,
      beam_node1, reference, node, 0.,ysc,zsc,
      beam_node2, reference, node, 0.,ysc,zsc,
      beam_node3, reference, node, 0.,ysc,zsc,
      eye,
      linear elastic generic, sym,
          EA, 0., 0., 0., z*EA, -y*EA,
          GAy, 0., 0., 0., 0.,
          GAz, 0., 0., 0., 0.,
          GJ, 0., 0., 0.,
          EJy+z^2*EA, -z*y*EA,

```

$$EJz+y^2*EA,$$

```
same,
same;
```

A piezoelectric actuator beam element is available; an arbitrary linear piezoelectric actuation matrix is required, together with the labels of the abstract nodes that represent the input signal tensions, as follows:

```
<normal_arglist> ::=
  <node_1_label> , (Vec3) <relative_offset_1> ,
  <node_2_label> , (Vec3) <relative_offset_2> ,
  <node_3_label> , (Vec3) <relative_offset_3> ,
  (OrientationMatrix) <orientation_matrix_section_I> ,
  (ConstitutiveLaw<6D>) <constitutive_law_section_I> ,
  { same | (OrientationMatrix) <orientation_matrix_section_II> } ,
  { same | (ConstitutiveLaw<6D>) <constitutive_law_section_II> } ,
  piezoelectric actuator ,
    <electrodes_number> ,
    <abstract_node_label_list> ,
    (Mat6xN) <piezoelectric_matrix_I> ,
    { same | (Mat6xN) <piezoelectric_matrix_II> }
  [ , <custom_output> ]
```

where the `abstract_node_label_list` is the list of the labels of the abstract nodes that represent the electrodes.

Private Data The following data is available:

1. "ex" axial strain
2. "ey" shear strain (local axis 2)
3. "ez" shear strain (local axis 3)
4. "kx" curvature about local axis 1 (torsional)
5. "ky" curvature about local axis 2 (bending)
6. "kz" curvature about local axis 3 (bending)
7. "Fx" axial force
8. "Fy" shear force (direction 2)
9. "Fz" shear force (direction 3)
10. "Mx" moment about local axis 1 (torsional)
11. "My" moment about local axis 2 (bending)
12. "Mz" moment about local axis 3 (bending)
13. "Xx" absolute position component 1
14. "Xy" absolute position component 2




Figure 8.5: Geometry of the two-node beam element.

15. "Xz" absolute position component 3
16. "Phix" absolute orientation vector component 1
17. "Phiy" absolute orientation vector component 2
18. "Phiz" absolute orientation vector component 3
19. "Omegax" absolute angular velocity component 1
20. "Omegay" absolute angular velocity component 2
21. "Omegaz" absolute angular velocity component 3
22. "ePx" axial strain rate
23. "ePy" shear strain rate (local axis 2)
24. "ePz" shear strain rate (local axis 3)
25. "kPx" curvature rate about local axis 1 (torsional)
26. "kPy" curvature rate about local axis 2 (bending)
27. "kPz" curvature rate about local axis 3 (bending)

Each string is prefixed by "pI." or "pII." to specify data related to either the first or the second point:

```
"pI.Fx"    # axial force at point I
"pII.kz"   # bending curvature about local axis 3 at point II
```

Private data related to point I are numbered from 1 to 27; private data related to point II are numbered from 28 to 54.

8.3.3 Two-node beam element

Similar considerations apply to the two-node beam. Its geometry is illustrated in Figure 8.5.

The syntax is

```

<elem_type> ::= beam2

<normal_arglist> ::=
  <node_1_label> , (Vec3) <relative_offset_1> ,
  <node_2_label> , (Vec3) <relative_offset_2> ,
  (OrientationMatrix) <orientation_matrix_section_I> ,
  (ConstitutiveLaw<6D>) <constitutive_law_section_I>
  [ , piezoelectric actuator ,
    <electrodes_number> ,
    <abstract_node_label_list> ,
    (Mat6xN) <piezoelectric_matrix_I> ]
  [ , <custom_output> ]

```

A more complete form is

```

<elem_type> ::= beam2

<normal_arglist> ::=
  <node_1_label> ,
  [ position , ] (Vec3) <relative_offset_1> ,
  [ orientation , (OrientationMatrix) <relative_orientation_1> , ]
  <node_2_label> ,
  [ position , ] (Vec3) <relative_offset_2> ,
  [ orientation , (OrientationMatrix) <relative_orientation_2> , ]
  { (OrientationMatrix) <orientation_matrix_section_I>
    | from nodes } ,
  (ConstitutiveLaw<6D>) <constitutive_law_section_I>
  [ , piezoelectric actuator ,
    <electrodes_number> ,
    <abstract_node_label_list> ,
    (Mat6xN) <piezoelectric_matrix_I> ]
  [ , <custom_output> ]

```

Private Data. The same private data indicated for the `beam3` element is available (see Section 8.3.2). No prefix must be specified ("`pI.`" is implicit).

Example. As an example, a simple beam element, with diagonal section stiffness matrix is presented:

```

set: integer beam_label = 1000;
set: integer beam_node1 = 2001;
set: integer beam_node2 = 2002;
set: real L = .4; # m
set: real EA = 1e6; # N
set: real GAy = .6e6; # N
set: real GAz = .6e6; # N
set: real GJ = 1.e3; # Nm^2
set: real EJy = 2.e3; # Nm^2
set: real EJz = 1.e4; # Nm^2
beam2: beam_label,
      beam_node1, reference, node, null,
      beam_node2, reference, node, null,

```

```

eye,
linear elastic generic, diag,
  EA, 1./(1./GAy+L^2/12./EJz), 1./(1./GAz+L^2/12./EJy),
  GJ, EJy, EJz;

```

Note that the shear terms have been naïvely inverted to eliminate shear locking, according to Equation (8.48).

8.3.4 Output

The output related to beam elements is contained in a file with extension `.act`; for each time step, the output is written for those element it was requested. The internal forces and moments are computed from the interpolated strains along the beam by means of the constitutive law, at the evaluation points. The format is:

- column 1: the label of the beam;
- columns 2–4: the three components of the force at the first evaluation point, oriented according to the reference frame of that beam section;
- columns 5–7: the three components of the moment at the first evaluation point, oriented according to the reference frame of that beam section.

The three-node beam element generates six additional columns:

- columns 8–10: the three components of the force at the second evaluation point, oriented according to the reference frame of that beam section;
- columns 11–13: the three components of the moment at the second evaluation point, oriented according to the reference frame of that beam section.

More detailed output is allowed when using NetCDF; see Section A.1.4 for details.

8.3.5 Notes

Two-node beam elements should be used with care.

8.4 Body

```

<elem_type> ::= body

<normal_arglist> ::= <node_label> ,
  { <one_body>
    | variable mass , <one_vm_body>
    | condense , <num_masses> , <one_body> [ , ... ] }

<one_body> ::=
  (Scalar) <mass> ,
  (Vec3) <relative_center_of_mass> ,
  (Mat3x3) <inertia_matrix>
  [ , inertial ,
    { node | (OrientationMatrix) <orientation_matrix> } ]

```

```

<one_vm_body> ::=
  (DriveCaller)      <mass> ,
  (TplDriveCaller<Vec3>) <relative_center_of_mass> ,
  (TplDriveCaller<Mat3x3>) <variable_mass_inertia_matrix> ,
  (TplDriveCaller<Mat3x3>) <variable_geometry_inertia_matrix>

```

If only one mass is defined, the first method should be used. Otherwise, many masses can be referred to the same element by means of the keyword `condense`, followed by the number of expected masses `num_masses`. The format of each sub-mass is the same as for the single mass input (actually, when `condense` is not supplied, `num_masses` is assumed to be 1).

The `inertia_matrix` is always referred to the center of mass of the mass that is being added. It can be rotated locally by means of the extra `orientation_matrix` supplied after the (optional) keyword `inertial`. The keyword `node` corresponds to the default, i.e. the inertia matrix is assumed to be input in the node reference frame.

Note: in many commercial finite element software, the off-diagonal elements of the inertia matrix are defined with a minus sign; for instance, NASTRAN's CONM2 lumped inertia card expects the values as indicated in Figure 8.6. However, the matrix is reconstructed as

$$\text{NASTRAN} ::= \begin{bmatrix} M & & & & & & & & \\ & M & & & & & & & \\ & & M & & & & & & \\ & & & \text{symmetric} & & & & & \\ & & & & I_{11} & & & & \\ & & & & -I_{21} & I_{22} & & & \\ & & & & -I_{31} & -I_{32} & I_{33} & & \end{bmatrix}$$

see for instance *NASTRAN V70.5 Quick Reference Guide* for details.

On the contrary, MBDyn directly reads the matrix that will be used in the computation, i.e. **without the minus signs in the off-diagonal terms**, as reported below:

$$\text{MBDyn} ::= \begin{bmatrix} i_{11} & i_{12} & i_{13} \\ & i_{22} & i_{23} \\ \text{sym.} & & i_{33} \end{bmatrix}$$

So:

$$\begin{aligned} i_{11} &= I_{11} \\ i_{22} &= I_{22} \\ i_{33} &= I_{33} \\ i_{12} &= -I_{21} \\ i_{13} &= -I_{31} \\ i_{23} &= -I_{32} \end{aligned}$$

The `variable mass` variant requires the mass, `mass`, the location of the center of mass with respect to the node in the reference frame of the node, `relative_center_of_mass`, and the inertia matrix about

\$2345678
CONM2	EID	G	CID	M	X1	X2	X3	
	I11	I21	I22	I31	I32	I33		

Figure 8.6: NASTRAN CONM2 card

the center of mass must be provided in form of *differentiable* drive callers of the specific dimension. The inertia matrix is logically split in two contributions: the first one, `variable_mass_inertia_matrix`, is related to the variation of mass; the second one, `variable_geometry_inertia_matrix`, is related to the change of geometry at fixed instantaneous value of the mass. The instantaneous inertia matrix with respect to the node, in the reference frame of the node, is computed as

$$\tilde{\mathbf{J}}_n = \text{variable_mass_inertia_matrix} \quad (8.49)$$

$$+ \text{variable_geometry_inertia_matrix} \quad (8.50)$$

$$+ \text{mass} \cdot \text{relative_center_of_mass} \times \text{relative_center_of_mass} \times^T. \quad (8.51)$$

See the technical manual for more details.

The inertia properties of the model can be logged and verified by means of the `inertia` keyword, as detailed in Section 8.17.3.

Private Data The following data is available:

1. **"E"** kinetic energy
2. **"V"** potential (i.e. gravitational) energy

Example.

```
set: integer NODE_LABEL = 100;
set: integer BODY_LABEL = 100;

# single mass example
body: BODY_LABEL, NODE_LABEL,
      8.,                               # mass
      reference, node, 0., 0., 0.,     # c.m. offset
      diag, 4.8, 4.8, .4;              # inertia tensor

# three masses example (equivalent to the previous one)
body: BODY_LABEL, NODE_LABEL,
      condense, 3,
      4.,                               # mass 1 (mid)
      reference, node, 0., 0., 0.,     # c.m. offset 1
      diag, .4, .4, .2,                # inertia tensor 1
      2.,                               # mass 2 (top)
      reference, node, 0., 0., 1.,     # c.m. offset 2
      diag, .2, .2, .1,                # inertia tensor 2
      2.,                               # mass 3 (bottom)
      reference, node, 0., 0., -1.,    # c.m. offset 3
      diag, .2, .2, .1;                # inertia tensor 3
```

8.5 Bulk Elements

The `bulk` element is intended as a sort of NASTRAN's CELAS card, that can be used to apply a stiffness term on an arbitrary degree of freedom. Extensions are planned to different kind of elements. The syntax of the `bulk` element is:

```
<elem_type> ::= bulk
```

```
<normal_arglist> ::= <bulk_type> , <bulk_arglist>
```

At present only the **stiffness spring** type is available.

8.5.1 Stiffness spring

```
<bulk_type> ::= stiffness spring
```

```
<bulk_arglist> ::=  
  (NodeDof) <dof> ,  
  (DriveCaller) <stiffness_drive>
```

The equation related to the desired dof of the linked node is added a contribution based on the value of the desired degree of freedom (even the derivative can be used) multiplied times the stiffness.

*Note: this family of elements has been partially superseded by the **genel** elements, which allow more generality.*

8.6 Couple

A variant of **force**; see Section 8.8 for details.

8.7 Electric Elements

electric elements are those elements that model electric and electronic devices, dealing with abstract degrees of freedom more than with electric ones (from the program's point of view they are exactly the same, the difference is only semantic). The true electric elements, such resistors, switches and so on, are classified as **electric bulk** elements. The syntax for **electric** elements is:

```
<normal_arglist> ::= <electric_type> , <electric_arglist>
```

The **electric** elements implemented at present are:

- **accelerometer**
- **displacement**
- **motor**
- **discrete control**

The syntax is described in the following.

8.7.1 Accelerometer

```
<electric_type> ::= accelerometer
```

```
<electric_arglist> ::=  
  { translational | rotational } ,  
  <struct_node_label> ,  
  <abstract_node_label> ,  
  ((Unit)Vec3) <measure_direction>  
  [ , position , (Vec3) <position> ]
```

The **position** is optional; it is meaningless for **rotational** accelerometers. The measure is taken along or about direction **measure_direction**; the vector is internally normalized to unity.

Legacy element: accelerometer with built-in transfer function

```
<electric_type> ::= accelerometer

<electric_arglist> ::= <struct_node_label> ,
  <abstract_node_label> ,
  ((Unit)Vec3) <measure_direction> ,
  (Scalar) <omega> ,
  (Scalar) <tau> ,
  (Scalar) <csi> ,
  (Scalar) <kappa>
```

The label **struct_node_label** defines the node whose acceleration is being measured; the label **abstract_node_label** defines the **abstract node** that will receive the output signal. An **electric node** can be used as well (?). The measure is taken along direction **measure_direction**; the vector is internally normalized to unity. The transfer function of the accelerometer is:

$$\frac{e_0}{a} = \text{kappa} \cdot \frac{\text{tau} \cdot s}{(1 + \text{tau} \cdot s)(1 + 2 \cdot \text{csi} \cdot (s/\text{omega}) + (s/\text{omega})^2)}$$

where e_0 is the output signal, a is the input (the acceleration) and s is Laplace's variable.

8.7.2 Displacement

```
<electric_type> ::= displacement

<electric_arglist> ::=
  <struct_node_1_label> , (Vec3) <relative_offset_1> ,
  <struct_node_2_label> , (Vec3) <relative_offset_2> ,
  <abstract_node_label>
```

$$\mathbf{d} = \mathbf{x}_b + \mathbf{o}_b - \mathbf{x}_a - \mathbf{o}_a \quad (8.52a)$$

$$x = \sqrt{\mathbf{d}^T \mathbf{d}} \quad (8.52b)$$

The value x is added to the right-hand side of the equation of the **abstract** node.

8.7.3 Motor

```
<electric_type> ::= motor

<electric_arglist> ::=
  <struct_node_1_label> ,
  <struct_node_2_label> ,
  ((Unit)Vec3) <direction_relative_to_node_1> ,
  <abstract_node_1_label> ,
  <abstract_node_2_label> ,
  (Scalar) <dG> ,
  (Scalar) <dL> ,
  (Scalar) <dR>
```

$$\boldsymbol{\omega} = \mathbf{e}^T (\boldsymbol{\omega}_2 - \boldsymbol{\omega}_1) \quad (8.53a)$$

$$V = V_2 - V_1 \quad (8.53b)$$

$$V = \mathbf{dL} \frac{di}{dt} + \mathbf{dR} i + \mathbf{dG} \boldsymbol{\omega} \quad (8.53c)$$

$$\mathbf{c}_1 = -\mathbf{e} \mathbf{dG} i \quad (8.53d)$$

$$\mathbf{c}_2 = \mathbf{e} \mathbf{dG} i \quad (8.53e)$$

$$i_1 = -i \quad (8.53f)$$

$$i_2 = i \quad (8.53g)$$

The motor applies an internal torque about direction $\mathbf{e} = \text{direction_relative_to_node_1}$ to both the structural nodes it is connected to. The direction \mathbf{e} is internally normalized to unity. It also applies a current to both the electrical nodes it is connected to.

The element assumes that the relative orientation between nodes 1 and 2 can only change by rotating about direction \mathbf{e} , so appropriate joints must be in place.

8.7.4 Discrete control

This element implements a discrete equation that can be used to represent the behavior of a discrete linear controller. The control matrices can be either provided, or identified during the simulation by recursive least squares.

Assume the original system consists in a generic nonlinear process that depends on a set of measures \mathbf{y} and a set of inputs \mathbf{u} at a finite number of time steps past the current time, t_k , namely

$$\mathbf{f}(\mathbf{y}_{k-i}, \mathbf{u}_{k-j}) = \mathbf{0}, \quad (8.54)$$

with $i = 0, p$ and $j = 0, q$. Only \mathbf{y}_k is unknown, and thus represents the output of the process. It is assumed that \mathbf{u}_k is known, and represents an input to the process.

This element implements a controller of the form

$$\mathbf{u}_{ck} = \sum_{i=1,p} \mathbf{B}_{ci} \mathbf{u}_{k-i} + \sum_{j=1,q} \mathbf{A}_{cj} \mathbf{y}_{k-j}, \quad (8.55)$$

where \mathbf{u}_{ck} is the control input that must be applied at time t_k in addition to any exogenous input \mathbf{u}_{ek} , so that $\mathbf{u}_k = \mathbf{u}_{ek} + \mathbf{u}_{ck}$. The control input \mathbf{u}_{ck} can only depend on the measures and the inputs at previous time steps.

Note that the symbols commonly used for discrete systems are here reversed, since the element is intended to compute the control signals at time t_k , \mathbf{u}_{ck} , based on the previous value of the controls, $\mathbf{u}_{k-i} = \mathbf{u}_{e(k-i)} + \mathbf{u}_{c(k-i)}$, and on the previous value of the motion of the original system, \mathbf{y}_{k-j} ; \mathbf{u}_e indicates a generic measured exogenous input to the system to be controlled.

The order p and q of the auto-regressive and exogenous portions of the system can differ. In detail, the order p can be zero; in that case, the system implements a *finite impulse response* function.

This element is not self-starting; it assumes that both inputs and outputs at times before the start time are zero.

The so-called “output” signals, indicated with \mathbf{y} , can be instances of either `NodeDof` or `DriveCaller` objects. The so-called “input” signals, indicated with \mathbf{u} , must be instances of `NodeDof` object. This implies that instances of `NodeDof` objects need the corresponding equations to be (at least) statically defined. In fact, the `discrete control` element uses the “output” `NodeDof` values, the \mathbf{y} , to compute the corresponding “inputs”, the \mathbf{u} . The latter are substantially added as a right-hand side to the equations of the corresponding `NodeDof` objects. As a consequence, other elements must contribute to the left-hand side of all the `NodeDof` equations in order to make them defined.

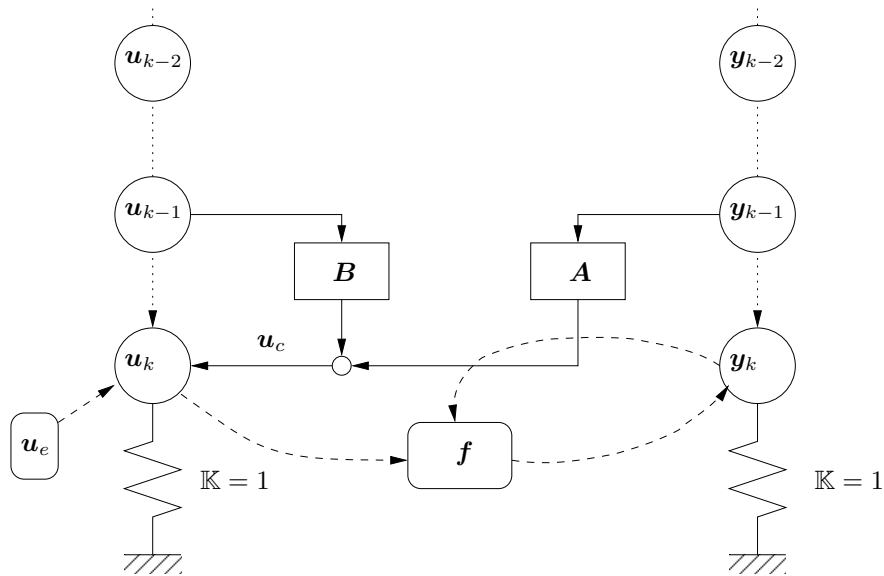


Figure 8.7: Discrete control layout.

Note that other elements may also contribute to the right-hand side of the “input” `NodeDof` object. Specifically, other `abstract` forces may contribute to their value. In that case, the additional forces represent the exogenous inputs. They are considered as part of the input used by the `discrete control` element, since the value u_k to be used in the control equation is extracted from the value of the corresponding `NodeDof` objects.

Figure 8.7 illustrates the behavior of the element. The typical suggested approach is illustrated later in an example.

The syntax is

```

<electric_type> ::= discrete control

<electric_arglist> ::= <num_outputs> , <num_inputs> ,
  <order_A> [ , fir , <order_B> ] ,
  <num_iter> ,
  <control_data> ,
  outputs ,
  (ScalarValue) <output_value>
  [ , scale , (DriveCaller) <scale> ]
  [ , ... ] ,
  inputs ,
  (NodeDof) <input_dof>
  [ , ... ]

<output_value> ::=
  { [ node dof , ] (NodeDof) <output_dof>
    | drive , (DriveCaller) <output_drive> }

```

The lists of the output and input dofs follows. The inputs do not require the `order_A` and `order_B` fields, since they are simply used to compute the control forces, and thus identify an equation. `order_B`

defaults to `order_A` unless a `fir` control is chosen.

The `control_data` has the syntax:

```
<control_data> ::= <control_type> , <control_arglist>
```

At present only a simple form of control is implemented. Other types to come are system identification, both recursive and one-shot, and adaptive control, with different models and schemes, all based on Generalized Predictive Control (GPC) and Deadbeat Control. The `control_data` syntax is:

```
<control_data> ::=  
  { control , " <control_matrices_file> "  
    | identification , <identification_data>  
    | adaptive control , <adaptive_control_data> }
```

Control

The file `control_matrices_file` must contain the matrices a_c , b_c of the control in plain text (as generated by Matlab, for instance):

```
 $a_{c1}$ ,  
...,  
 $a_{cp}$ ,  
 $b_{c1}$ ,  
...,  
 $b_{cp}$ 
```

where p is the `order_A` of the controller and the matrices a_c have `num_inputs` rows and `num_outputs` columns, while the matrices b_c have `num_inputs` rows and `num_inputs` columns.

Identification

```
<identification_data> ::=  
  { arx | armax }  
  [ , <forgetting_factor> ]  
  [ , <persistent_excitation> ]  
  [ , file , " <output_file_name> " ]
```

The forgetting factor is defined as

```
<forgetting_factor> ::=  
  forgettingfactor ,  
  { const , <d>  
    | dynamic , <n1> , <n2> , <rho> , <fact> , <kref> , <klim> }
```

The default is a `const` forgetting factor with $d = 1$.

The `persistent_excitation` is defined as

```
<persistent_excitation> ::=  
  excitation , (DriveCaller) <excitation_drive> [ , ... ]
```

where `num_inputs` instances of `excitation_drive` must be defined. By default, no persistent excitation is defined.

Adaptive control

The `adaptive_control_data` card is

```
<adaptive_control_data> ::=
  [ { arx | armax } , ]
  [ periodic , <periodic_factor> , ]
  { gpc ,
    <prediction_advancing_horizon> ,
    <control_advancing_horizon> ,
    <prediction_receding_horizon> ,
    [ prediction_weights , <Wi> [ , ... ] , ]
    [ control_weights , <Ri> [ , ... ] , ]
    (DriveCaller) <weight_drive>
  | deadbeat ,
    <prediction_advancing_horizon> ,
    <control_advancing_horizon> }
  [ , <forgetting_factor> ]
  [ , <persistent_excitation> ]
  [ , trigger , (DriveCaller) <trigger_drive> ]
  [ , desired_output , (DriveCaller) <output_drive> [ , ... ] ]
  [ , file , " <file_name> " ]
```

The default is `arx`.

The meaning of the keywords is:

- the `periodic_factor` defaults to 0;
- if the keyword `prediction_weights` is present, a number of weights `Wi` equal to

$$(\text{prediction_advancing_horizon} - \text{prediction_receding_horizon})$$

must be supplied. If the keyword `control_weights` is present, a number of weights `Ri` equal to `control_advancing_horizon` must be supplied. If the keywords are not defined, the corresponding weights default to 1;

- the `desired_output` option requires `num_outputs` drives to be defined.

Private Data

The following data is available:

1. `"u[<idx>]"` value of the `idx`-th control output.

Examples

Consider the case of a discrete controller that computes a set of control signals \mathbf{u}_k by adding to their value at the previous step, \mathbf{u}_{k-1} , a contribution coming from some measure \mathbf{y}_{k-1} . The control matrices are

$$\mathbf{A}_{c1} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad (8.56a)$$

$$\mathbf{B}_{c1} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (8.56b)$$

They are defined in the control data file `discretecontrol.dat` as

```

a_11 a_12
a_21 a_22
1.0 0.0
0.0 1.0

```

The model consists in two abstract nodes for the control input `NodeDof` objects. They need to be grounded. This can be formulated by using a `spring support genel` for each `abstract` node, characterized by a unit spring coefficient, such that

$$1 \cdot x_{u1} = u_1 \quad (8.57a)$$

$$1 \cdot x_{u2} = u_2 \quad (8.57b)$$

Moreover, assume that the first measure, y_1 , comes from a `NodeDof` object, while the second one, y_2 , comes from a `DriveCaller` object. The equation corresponding to y_1 must be appropriately grounded as well, for example by means of yet another `spring support genel`. This is outside the scope of the present example, so it will assume the output nodes are defined appropriately.

The model is

```

set: integer U_1 = 1;
set: integer U_2 = 2;
set: integer Y_1 = 11;
# ...
abstract: U_1;
abstract: U_2;
# ...
genel: U_1, spring support,
      U_1, abstract, algebraic,
      linear elastic, 1.;
genel: U_2, spring support,
      U_2, abstract, algebraic,
      linear elastic, 1.;
electric: 99, discrete control,
         2, # number of 'outputs' y
         2, # number of 'inputs' u
         1, # order of output matrices A (same for input matrices B)
         1, # update every iteration
         control, "discretecontrol.dat",
         outputs,
           node dof, Y_1, abstract, algebraic, scale, 1.e-3,
           drive, sine, .5, 2*pi, 1., forever, 0, scale, 1.e-3,
         inputs,
           U_1, abstract, algebraic,
           U_2, abstract, algebraic;

```

8.8 Force

The `force` element is a general means to introduce a right-hand side to the equations, i.e. an explicit contribution to the equations. There is a basic distinction between abstract and structural forces: abstract forces apply to arbitrary equations, while structural forces (and couples) are specific to structural nodes and have a spatial characterization. Essentially, structural forces have three components that may

depend on arbitrary parameters (usually the simulated time), and a location in space. Structural couples have three parameter-dependent components. The syntax of the `force` element is:

```

<elem_type> ::= { force | couple }

<normal_arglist> ::= <force_type> , <force_arglist>

<force_type> ::=
  { abstract [ internal ]                # force
    | { absolute | follower | total } [ internal ] # force, couple
    | modal                               # force
    | external { structural | modal } [ mapping ] } # force

```

The `abstract force_type` applies to generic equations;

The structural force types (`absolute`, `follower`, `total`) only apply to equilibrium equations related to `structural` nodes. The `couple` element can only be structural. It is discussed in this section because of its input syntax commonality with the structural `force` elements.

The `modal force_type` applies to equations related to a `modal` element.

The `external structural` and the `external modal` types are essentially used to provide an easy to use interface for coupling with external software in a loose or tight manner.

8.8.1 Output

The output is discussed according to the types of forces. The label of the element is output first in all the cases.

8.8.2 Abstract force

```

<elem_type> ::= force

<force_type> ::= abstract

<force_arglist> ::=
  (NodeDof) <dof> ,
  (DriveCaller) <force_magnitude>

```

the `dof` field is a normal `NodeDof`; no `order` is required since the `force` simply applies to the equation related to the node, regardless of the order.

Output

The format is:

- the label of the element;
- the label of the node¹ the force is applied to;
- the value of the force.

¹Since the `abstract` force type can connect to `NodeDof`, information provided here can be partial if the `NodeDof` actually represents a component of a node with more than one degree of freedom.

8.8.3 Abstract reaction force

```
<elem_type> ::= force

<force_type> ::= abstract internal

<force_arglist> ::=
  (NodeDof)    <dof_1> ,
  (NodeDof)    <dof_2> ,
  (DriveCaller) <force_magnitude>
```

the `dof_1` and `dof_2` fields are normal `NodeDof` but no `order` is required since the `force` simply applies to the equations related to the nodes, regardless of the order, with opposite magnitudes.

Output

The format is:

- the label of the element;
- the label of the first node² the force is applied to;
- the value of the force applied to the first node;
- the label of the second node³ the force is applied to;
- the value of the force applied to the second node (opposite to that applied to the first node).

8.8.4 Structural force

```
<elem_type> ::= force

<force_type> ::= { absolute | follower }

<force_arglist> ::=
  <node_label> ,
  position , (Vec3) <relative_arm> ,
  (TplDriveCaller<Vec3>) <force_value>

<force_type> ::= total

<force_arglist> ::=
  <node_label>
  [ , position , (Vec3) <relative_arm> ]
  [ , force orientation , (Mat3x3) <force_orientation> ]
  [ , moment orientation , (Mat3x3) <moment_orientation> ]
  [ , force , (TplDriveCaller<Vec3>) <force_value> ]
  [ , moment , (TplDriveCaller<Vec3>) <moment_value> ]
```

²See footnote 1.

³See footnote 1.

The vector `relative_arm` defines the offset with respect to the node of the point where the force is applied. The drive `force_value` defines the value of the force as a function of time. The `total` force is intrinsically follower. See also the `total joint` (Section 8.12.46).

The force and the moment applied to node `node_label` are thus

- `absolute` force element:

$$\mathbf{f} = \text{force_value} \quad (8.58a)$$

$$\mathbf{m} = (\mathbf{R} \cdot \text{relative_arm}) \times \text{force_value}; \quad (8.58b)$$

- `follower` force element:

$$\mathbf{f} = \mathbf{R} \text{ force_value} \quad (8.59a)$$

$$\mathbf{m} = \mathbf{R}(\text{relative_arm} \times \text{force_value}); \quad (8.59b)$$

- `total` force element:

$$\mathbf{f} = \mathbf{R} \text{ force_orientation} \cdot \text{force_value} \quad (8.60a)$$

$$\begin{aligned} \mathbf{m} = \mathbf{R}(\text{relative_arm} \times (\text{force_orientation} \cdot \text{force_value}) \\ + \text{moment_orientation} \cdot \text{moment_value}); \end{aligned} \quad (8.60b)$$

\mathbf{R} is the orientation matrix of the node `node_label`.

Example.

```
force: 10, absolute,
  1000,
  position, .5, 0., 0.,
  1., 0., 0.,
  const, 25.;
force: 20, follower,
  2000,
  position, null,
  component,
  const, 0.,
  ramp, 10., 0., 1., 0.,
  file, 5, 2;
force: 30, total,
  2000,
  position, .5, 0., 0.,
  force orientation, eye,
  moment orientation, eye,
  1., 0., 0.,
  const, 25.,
  component,
  const, 0.,
  ramp, 10., 0., 1., 0.,
  file, 5, 2;
```

Output

The format is:

- the label of the element;
- the label of the node the force is applied to;
- the three components of the force;
- the arm of the force, in the global frame (i.e. referred to point $\{0,0,0\}$ and oriented as the global frame)

8.8.5 Structural internal force

```
<elem_type> ::= force

<force_type> ::= { absolute | follower } internal

<force_arglist> ::=
  <node_1_label> ,
  position , (Vec3) <relative_arm_1> ,
  <node_2_label> ,
  position , (Vec3) <relative_arm_2> ,
  (TplDriveCaller<Vec3>) <force_value>

<force_type> ::= total internal

<force_arglist> ::=
  <node_1_label> ,
  [ position , (Vec3) <relative_arm_1> , ]
  [ force orientation , (Mat3x3) <force_orientation_1> , ]
  [ moment orientation , (Mat3x3) <moment_orientation_1> , ]
  <node_2_label>
  [ , position , (Vec3) <relative_arm_2> ]
  [ , force orientation , (Mat3x3) <force_orientation_2> ]
  [ , moment orientation , (Mat3x3) <moment_orientation_2> ]
  [ , force , (TplDriveCaller<Vec3>) <force_value> ]
  [ , moment , (TplDriveCaller<Vec3>) <moment_value> ]
```

The format is basically identical to the previous case, except for the definition of the second node.

Output

The format is:

- the label of the element;
- the label of the first node the force is applied to;
- the three components of the force applied to the first node;
- the arm of the force with respect to the first node, in the global frame (i.e. referred to point $\{0,0,0\}$ and oriented as the global frame).

- the label of the second node the force is applied to;
- the three components of the force applied to the second node (opposite to that applied to the first node);
- the arm of the force with respect to the second node, in the global frame (i.e. referred to point $\{0,0,0\}$ and oriented as the global frame).

Example.

```
# constant structural force
force: 1, absolute,
      10, position, null,
      0., 0., 1.,
      const, 100.;

# constant structural internal force
force: 1, absolute internal,
      10, position, null,
      20, position, null,
      0.,0.,1.,
      const, 100.;
```

8.8.6 Structural couple

The structural couple is defined by using `couple` as `elem_type`.

```
<elem_type> ::= couple

<force_type> ::= { absolute | follower }

<force_arglist> ::=
  <node_label> ,
  [ position , (Vec3) <relative_arm> , ]
  (TplDriveCaller<Vec3>) <couple_value>
```

The vector `relative_arm` defines the offset with respect to the node of the point where the couple is applied. It is not used in the analysis (since it makes no sense), but it can be optionally provided for future reference (e.g. for the visualization of the couple). The drive `couple_value` defines the value of the couple as a function of time.

The moment applied to node `node_label` is thus

- `absolute` couple element:

$$m = \text{couple_value} \quad (8.61)$$

- `follower` couple element:

$$m = R \cdot \text{couple_value} \quad (8.62)$$

R is the orientation matrix of the node `node_label`.

Output

The format is:

- the label of the element;
- the label of the node the couple is applied to;
- the three components of the couple.

8.8.7 Structural internal couple

The structural couple is defined by using `couple` as `elem_type`.

```
<elem_type> ::= couple

<force_type> ::= { absolute | follower } internal

<force_arglist> ::=
  <node_1_label> ,
  [ position , (Vec3) <relative_arm_1> , ]
  <node_2_label> ,
  [ position , (Vec3) <relative_arm_2> , ]
  (TplDriveCaller<Vec3>) <couple_value>
```

The format is basically identical to the previous case, except for the definition of the second node.

Output

The format is:

- the label of the element;
- the label of the first node the couple is applied to;
- the three components of the couple applied to the first node;
- the label of the second node the couple is applied to;
- the three components of the couple applied to the second node (opposite to that applied to the first node);

Note: by using a `dof`, a `node` or an `element` drive, a simple feedback control can be easily implemented. Note however that the dependence of a force on some internal state does not result in adding the corresponding contribution to the Jacobian matrix: the force remains explicit. Their improper use may result in missing convergence, or singularity of the Jacobian matrix of the problem.

8.8.8 Modal

This element allows to define a set of generalized forces acting on the equations related to the generalized variables of a `modal` joint element.

```

<elem_type> ::= force

<force_type> ::= modal

<force_arglist> ::=
  <modal_label> ,
  [ list , <number_of_modes> , <mode1> [ , ... ] , ]
  (DriveCaller) <force1>
  [ , resultant , (Vec3) <f> , (Vec3) <m> ]
  [ , ... ]

```

It requires the label of the modal joint it is connected to, `modal_label`. If the optional keyword `list` is given, the number of excited modes `number_of_modes` and the list of the excited modes is expected, otherwise it assumes all modes of the `modal_label` modal joint will be excited.

A list of `drive callers` (see Section 2.5). that provide the excitation value for each mode is required last. If the keyword `list` was given, the drive callers must be `number_of_modes`, otherwise they must be as many as the modes of the excited modal joint element.

If the `modal` joint element is connected to a modal node instead of being clamped to the ground, the optional keyword `resultant` is allowed for each mode drive caller. The keyword `resultant` is followed by two 3×1 vectors containing the resultant force and moment associated to the related mode.

If the keyword `list` was given, the mode numbers must be the actual mode numbers defined in the FEM model.

Example. In this example, all the modes available in the FEM data file are used by the modal joint, and all the modes are excited by the modal force

```

joint: MODAL_JOINT, modal, MODAL_NODE,
      5,           # number of modes
      from file,
      "model.fem",
      0;          # no interface nodes
force: MODAL_FORCE, modal, MODAL_JOINT,
      const, 5.,
      const, 10.,
      const, 1.,
      const, -1.,
      const, -3.;

```

In this example, only three modes of the same FEM data file of the previous example are used by the modal joint; only the last mode is excited by the modal force

```

joint: MODAL_JOINT, modal, MODAL_NODE,
      3,           # number of modes
      list, 1, 3, 5,
      from file,
      "model.fem",
      0;          # no interface nodes
force: MODAL_FORCE, modal, MODAL_JOINT,
      list, 1,    # only one mode is excited
      5,         # it is mode #5 in "model.fem"
      const, -3.;

```

Note that the mode is indicated by its number in the FEM data file.

8.8.9 External forces

External forces are elements that allow to communicate with an external software that computes forces based on information on the kinematics of the model.

Different communication schemes are supported. By default, communication occurs via files. Currently supported communication schemes are

```
<external_force_communicator> ::= { <file> | <edge> | <socket> }

<file> ::=
    # the default
    " <input_file_name> " [ , unlink ] ,
    " <output_file_name> " [ , no clobber ]
    [ , <common_parameters> ]

<edge> ::= edge ,
    " <flag_file_name> " ,
    " <data_file_name> "
    [ , <common_parameters> ]

<socket> ::= socket ,
    [ create , { yes | no } , ]
    { path , <path> | port , <port> [ , host , <host> ] }
    [ , <common_parameters> ]

<common_parameters> ::= <common_parameter> [ , ... ]

<common_parameter> ::=
    { sleep time , <sleep_time>
      | precision , { default | <digits> } # not valid for socket
      | coupling , { staggered | loose | tight | <coupling_steps> }
      | send after predict , { yes | no } }
```

File communicator

`input_file_name` is the name of the file MBDyn expects to find with the values of the force. The content depends on the purpose it is used for. The optional keyword `unlink` indicates that MBDyn is supposed to unlink the file as soon as it is ready to read another one, as a primitive form of inter-process communication.

`output_file_name` is the name of the file MBDyn creates each time it intends to communicate the kinematics of the model to the external software. The contents depend on the purpose it is used for. The option `no clobber` indicates that MBDyn is supposed to wait until the external software removed the file before generating a new one, as a primitive form of inter-process communication.

EDGE communicator

This communicator has been developed in cooperation with Luca Cavagna, to support communication with FOI's EDGE.

The file `flag_file_name` is used to synchronize the communication between the processes. The same file is truncated and reused by both processes. The syntax of the file changed since EDGE 5.1; this change is reflected from MBDyn 1.3.16 on. The file contains a two lines header, followed by a line containing a single digit. The header is

```
UPDATE,N,0,0,1
FLAG,I,1,1,0
```

The digit indicates the command. Its meaning is:

- 0 : external software is initializing; MBDyn must wait;
- 1 : external software is busy; MBDyn must wait;
- 2 : external software is ready to read kinematics; MBDyn can start reading forces and writing kinematics;
- 3 : MBDyn finished writing kinematics; external software can start reading kinematics and writing forces (written by MBDyn when data file is ready);
- 4 : external software converged; MBDyn should go to convergence, advance to the next time step, and start writing new step kinematics;
- 5 : external software decided to quit; MBDyn should stop communicating.

So far, there are no provisions for re-executing a time step, nor for communicating either the time or the time step, so adaptive time stepping cannot be implemented.

`data_file_name` is the name of the file used to actually send data between the two processes. Its contents depend on the purpose it is used for, as described in the subsequent sections.

Socket communicator

The optional parameter `create`, when set to `yes`, indicates that MBDyn will create the socket, and the peer will have to connect to it. Otherwise, when set to `no`, it indicates that MBDyn will try to connect to an already existing socket created by the peer. Connecting to a peer is attempted while reading the input file. Sockets are created when the input file has been read. MBDyn waits until all sockets have been connected to by peers before the simulation is started.

MBDyn supports local (unix) sockets, defined using the `path` parameter, and inet sockets, defined using the `port` parameter. When `create` is set to `yes`, the optional keyword `host` allows to define what interface MBDyn will listen on. When `create` is set to `no`, the optional keyword `host` indicates what host to connect to. It defaults to 'localhost'.

Common parameters

The optional parameter `sleep_time` determines how long MBDyn is supposed to sleep while waiting for a new input file to appear or for an old output file to disappear.

The optional parameter `precision` determines the precision used in writing the output file; the default is 16 digits. This is only meaningful for communicators based on textual files.

The optional parameter `coupling` determines whether the coupling will be `staggered` (communicate each time step, with one step delay), `loose` (communicate each time step, no delay, the default) or `tight` (communicate each iteration); otherwise, if `coupling_steps` is provided, the communication occurs every `coupling_steps`.

The optional parameter `send after predict` allows to indicate whether MBDyn must send the predicted motion or not when playing a tight coupling loop.

Staggered Coupling. When the coupling is **staggered**, the communication pattern is:

1. MBDyn receives a set of forces sent by the external peer the first time the residual is assembled within time step k ; the external peer generated that set of forces based on the kinematics at step $k - 1$;
2. MBDyn sends a set of kinematic parameters related to step k after convergence.

As a consequence, MBDyn solves the kinematics at the current time step, k , using forces evaluated for the kinematics at the previous time step, $k - 1$. Something like

$$\mathbf{f}(\mathbf{y}_k, \dot{\mathbf{y}}_k, \mathbf{u}_{k-1}) = \mathbf{0}. \quad (8.63)$$

Note: not implemented yet.

Loose Coupling. When the coupling is **loose**, the communication pattern is:

1. MBDyn sends a set of kinematic parameters related to step k after predict.
2. MBDyn receives a set of forces sent by the external peer the first time the residual is assembled within time step k ; the external peer generated that set of forces based on the predicted kinematics at step k ;

As a consequence, MBDyn solves the kinematics at the current time step, k , using forces evaluated for the predicted kinematics at the current time step, k . Something like

$$\mathbf{f}(\mathbf{y}_k, \dot{\mathbf{y}}_k, \mathbf{u}_k^{(0)}) = \mathbf{0}. \quad (8.64)$$

Tight Coupling. When the coupling is **tight**, the communication pattern is:

1. MBDyn sends the predicted kinematics for step k during the “after predict” phase (unless **send after predict** is set to **no**);
2. MBDyn receives a set of forces sent by the external peer each time the residual is assembled; those forces are computed based on the kinematics at iteration j
3. as soon as, while reading the forces, MBDyn is informed that the external peer converged, it continues iterating until convergence using the last set of forces it read.

As a consequence, MBDyn solves the kinematics at the current time step, k , at iteration j , using forces evaluated for the kinematics at the same time step, k , but at iteration $j - 1$. Something like

$$\mathbf{f}(\mathbf{y}_k^{(j)}, \dot{\mathbf{y}}_k^{(j)}, \mathbf{u}_k^{(j-1)}) = \mathbf{0}. \quad (8.65)$$

8.8.10 External Structural

This element allows to communicate with an external software that computes forces applied to a pool of nodes and may depend on the kinematics of those nodes. Communication occurs by means of the communicators illustrated earlier.

```
<elem_type> ::= force
<force_type> ::= external structural
```

```

<force_arglist> ::= <external_force_communicator> ,
  [ reference node , <ref_node_label> , ]
  [ { labels , { yes | no }
    | sorted , { yes | no }
    | orientation ,
      { none | orientation matrix | orientation vector | euler 123 }
    | accelerations , { yes | no }
    | use reference node forces , { yes | no }
      [ , rotate reference node forces , { yes | no } ] ]
  [ , ... ] , ]
<num_nodes> ,
  <node_label> [ , offset , (Vec3) <offset> ]
  [ , ... ]
  [ , echo , <echo_file_name> ]

```

- **use reference node forces** only meaningful when **reference node** is used. It assumes the external solver is sending the forces and moments related to the reference node. They correspond to the rigid-body forces and moments applied to the whole system. As such, the forces and moments applied to each node are removed accordingly. If it is set to **no**, reference node forces and moments will be ignored.
- The orientation style **none** implies that only positions, velocities and accelerations will be output (the latter only if **accelerations** is set to **yes**).
- **echo** causes the output of the reference configuration to be produced in a file called **echo_file_name**. If the file exists, it is overwritten, if allowed by file system permissions. The format is that of the communicator in stream form.

File communicator

The forces are read from the input file in textual form, each line formatted as follows:

- a label (if **labels** is set to **yes**);
- three components of force in the global frame;
- three components of moment in the global frame, referred to the node as the pole.

The label indicates what node the force and the moment apply to; each force is applied in a point that may be optionally offset from the corresponding node location according to **offset**. If optional keyword **sorted** is set to **no**, the forces might be read in arbitrary order, so they need to be recognized by the label. The option **sorted** is only meaningful when **labels** is set to **yes**.

The kinematics are written to the output file in textual form, each line formatted as follows:

- a label (if **labels** is set to **yes**);
- the position in the global frame;
- the orientation of the node with respect to the global frame:
 - if **orientation** is set to **orientation matrix**, the orientation matrix in the global frame, row-oriented: $R_{11}, R_{12}, R_{13}, R_{21}, \dots, R_{32}, R_{33}$;
 - if **orientation** is set to **orientation vector**, the orientation vector's components;

- if **orientation** is set to **euler 123**, the three Euler angles, in degrees, according to the 1,2,3 convention;
- the velocity with respect to the global frame;
- the angular velocity with respect to the global frame.

If the optional keyword **accelerations** is set to **yes**, at the end of each line the following optional fields will appear:

- the acceleration with respect to the global frame;
- the angular acceleration with respect to the global frame.

The label indicates what node the kinematics is related to; the position and the velocity refer to a point that may be optionally offset from the node location according to **offset**.

If a **reference node** was given, the first record is related to the reference node itself. The following records contain forces and moments in the reference frame of the reference node.

EDGE communicator

TBD

Socket communicator

TBD

Example. The following defines an **external structural** force with the file communicator

```
set: integer FORCE_LABEL = 1000;
set: integer NODE_1 = 10;
set: integer NODE_2 = 20;
force: FORCE_LABEL, external structural,
      "IN.dat", unlink,
      "OUT.dat", no clobber,
      sleep time, 10,
      precision, 12,
      coupling, loose,
      2, # number of nodes
      NODE_1,
      NODE_2, offset, 0.,0.,.5;
```

Output

An external structural element writes one line for each connected node at each time step in the **.frc** file. Each line contains:

- in column 1 the label of the element and that of the corresponding node; the format of this field is **element_label@node_label**;
- in columns 2–4 the three components of the force;
- in columns 5–7 the three components of the moment.

If a reference node is defined, a special line is output for the reference node, containing:

- in column 1 the label of the element and that of the corresponding node; the format of this field is `element_label#node_label`;
- in columns 2–4 the three components of the force applied to the reference node, as received from the peer;
- in columns 5–7 the three components of the moment applied to the reference node, as received from the peer.
- in columns 8–10 the three components of the force that are actually applied to the reference node, in the global reference frame;
- in columns 11–13 the three components of the moment with respect to the reference node that are actually applied to the reference node, in the global reference frame;
- in columns 14–16 the three components of the force resulting from the combination of all nodal forces, in the global reference frame;
- in columns 17–19 the three components of the moment with respect to the reference node resulting from the combination of all nodal forces and moments, in the global reference frame;

8.8.11 External structural mapping

This element allows to communicate with an external software that computes forces applied to a pool of nodes and may depend on the kinematics of those nodes through a linear mapping. Communication occurs by means of the communicators illustrated earlier.

```

<elem_type> ::= force

<force_type> ::= external structural mapping

<force_arglist> ::= <external_force_communicator> ,
  [ reference node , <ref_node_label> , ]
  [ { labels , { yes | no }
    | orientation ,
      { none | orientation matrix | orientation vector | euler 123 }
    | accelerations , { yes | no }
    | use reference node forces , { yes | no }
      [ , rotate reference node forces , { yes | no } ] ]
  [ , ... ] , ]
<num_points> ,
  <node_label> , offset [ , <point_label> ] , (Vec3) <offset>
  [ , ... ]
[ , echo , " <echo_file_name> "
  [ , { precision , <digits>
    | surface , " <surface_file_name> "
    | output , " <output_file_name> "
    | order , <order>
    | basenode , <basenode>
    | weight , { inf | (Scalar) <weight> } }
  [ , ... ] ]
[ , stop ] ]

```

```
[ , mapped points number , { from file | <mapped_points> } ,
  { full | sparse } mapping file , " <mapping_file_name> "
  [ , threshold , <threshold> ]
  [ , { mapped labels file , " <mapped_labels_file_name> "
    | <mapped_label> [ , ... ] } ] ]
```

A set of `num_points` points is generated first, by computing the kinematics of the points originating from the rigid-body motion of the structural nodes `node_label` according to the specified offset. Multiple `offset` blocks can appear for each `node_label`.

The `echo` keyword causes the positions of those points to be written in a file named `echo_file_name`. If the keyword `stop` is present, the simulation ends here. This is useful to generate the bulk data that is needed to compute the linear mapping matrix. The optional keyword `precision` allows to set `precision` digits in the logged data; the other optional keywords refer to values that are written verbatim to the file. In detail, `surface_file_name` is the name of the file that contains the coordinates of the peer's points, `output_file_name` is the name of the file that will contain the mapping matrix, `order` is an integer ($1 \leq \text{order} \leq 3$), `basenode` is the number of nodes for each region (`basenode` > 0), `weight` is a weighting coefficient (`weight` ≥ -2 or `inf`)

When the `reference node` keyword is present, the kinematics of the points is formulated in the reference frame of node `ref_node_label`. The forces are expected in the same reference frame. In this case, if `use reference node forces` is set to `yes`, the force and moment related to node `ref_node_label` returned by the external solver are used; if `rotate reference node forces` is set to `yes`, they are first rotated in the global reference frame.

When the `mapped points number` keyword is present, a mapping matrix is read from the file named `mapping_file_name` (same as `output_file_name` when the `echo` keyword is used). If the keyword `from file` is used, the number of `mapped_points` is computed from the number of rows of the matrix in file `mapping_file_name`. The matrix must be $(3 \times \text{mapped_points}) \times (3 \times \text{num_points})$. If the `mapped points number` keyword is not present, the element behaves as if the mapping were the identity matrix, namely the position and the velocity of the points rigidly offset from MBDyn's nodes are directly passed to the peer solver. See the `modal` variant for a description of the mapping file format.

The labels of the original points, indicated as `point_label`, and the labels of the mapped points, either contained in the `mapped_labels_file_name`, or directly listed as `mapped_label`, must be provided when `labels` is set to `yes`.

Mapping. The mapping consists in a constant matrix that allows to compute the position and the velocity of the mapped points (subscript 'peer') as functions of a set of points rigidly offset from MBDyn's nodes (subscript 'mbdyn'),

$$\mathbf{x}_{\text{peer}} = \mathbf{H} \mathbf{x}_{\text{mbdyn}} \quad (8.66a)$$

$$\dot{\mathbf{x}}_{\text{peer}} = \mathbf{H} \dot{\mathbf{x}}_{\text{mbdyn}} \quad (8.66b)$$

The same matrix is used to map back the forces onto MBDyn's nodes based on the preservation of the work done in the two domains,

$$\delta \mathbf{x}_{\text{mbdyn}}^T \mathbf{f}_{\text{mbdyn}} = \delta \mathbf{x}_{\text{peer}}^T \mathbf{f}_{\text{peer}} = \delta \mathbf{x}_{\text{mbdyn}}^T \mathbf{H}^T \mathbf{f}_{\text{peer}} \quad (8.67)$$

which implies

$$\mathbf{f}_{\text{mbdyn}} = \mathbf{H}^T \mathbf{f}_{\text{peer}} \quad (8.68)$$

To compute matrix \mathbf{H} one needs to:

- gather the coordinates of the set of points used by the peer process (subscript ‘peer’ in the formulas above) in a file, formatted as

```
<num_points>
<x_1> <y_1> <z_1>
...
<x_n> <y_n> <z_n>
```

If **reference node** is defined, the points need to be expressed in a reference frame coincident with that of node **ref_node_label**.

- execute MBDyn’s input file after uncommenting the **echo** line in the definition of the external structural mapping force element, terminated by the **stop** keyword, to generate the **echo_file_name** file with MBDyn’s points coordinates (subscript ‘mbdyn’ in the formulas above).

- start octave

- make sure the folder **contrib/MLS/** is in octave’s path

```
addpath('/path/to/contrib/MLS')
```

- execute the **create_mls_interface** function

```
create_mls_interface('echo_file_name');
```

This generates the mapping matrix H , using the files **echo_file_name** and **surface_file_name** (the latter is written in **echo_file_name** by the **surface** keyword). The matrix is stored in the file **output_file_name**. The **sparse** storage is recommended, since for usual problems the matrix is significantly sparse.

NOTE: this operation can take minutes when the mapped domains consist of thousands of points.

- comment the **echo** line in the definition of the external structural mapping force element.

From this point on, when MBDyn is executed it loads the mapping matrix from file **output_file_name** using sparse storage, automatically ignoring coefficients whose absolute value is below **threshold** (defaults to 0).

The parameters of the mapping, as documented in **create_mls_interface**, are

```
% function create_mls_interface([inputfilename])
%     inputfilename [optional argument] string that contains the name
%     of the file created by the MBDyn external mapping force element
%     using the "echo" option
%     Additional rows must be either added in the comment section of the file
%     or can be enforced using the appropriate options of the "echo" keyword
%     if input values different from default are required. In detail
%
%     # surface: ...
%         used to indicate the name of the file that contains
%         the external element grid data. Default: 'surface.dat'
%
%     # order: used to indicate the polynomial order of the base used by MLS
%         The value must be in the range [1,3]. Default is 2
%
%     # basenode: number of nodes included in the local MLS support
```

```

%           should be higher than 0
%
%   # weight: continuity order of the RBF weight function
%           The value must be in the range [-2,inf]
%           -1 means constant weight == 1 in the support
%           -2 means a special weight function == 0 @ the query point
%           default 2
%
%   # output: name of the file where the interface matrix is saved in sparse format

```

Example.

```

force: 100, external structural mapping,
socket,
    create, yes,
    path, "/tmp/mbdyn.sock",
    no signal,
coupling, tight,
reference node, 0,
points number, 9,
    10,
        offset, null,          # centered at node
        offset, 0., 1., 0., # y offset
        offset, 0., 0., 1., # z offset
    20,
        offset, null,          # ...
        offset, 0., 1., 0.,
        offset, 0., 0., 1.,
    30,
        offset, null,
        offset, 0., 1., 0.,
        offset, 0., 0., 1.,
# NOTE: uncomment the line below to generate data required for the mapping
# echo, "points.dat", surface, "surf.dat", output, "mapping.dat", stop,
mapped points number, 27,
sparse mapping file, "mapping.dat";

```

File communicator

TBD

EDGE communicator

TBD

Socket communicator

TBD

The format of the socket communicator is relatively simple; however, implementors may exploit the peer library `libmbc`, whose C API is defined in `mbc.h`. A C++ wrapper is defined in `mbcxx.h`. A python wrapper is also available.

Output

TBD

8.8.12 External modal

This element allows to communicate with an external software that computes forces applied to a **modal** joint element, and may depend on the values of the corresponding generalized variables. Communication occurs by means of the communicators illustrated earlier.

```
<elem_type> ::= force

<force_type> ::= external modal

<force_arglist> ::= <external_force_communicator> ,
    <modal_label>
    [ , accelerations ]
    [ , type , { rigid | modal | all } ]
```

modal_label is the label of the **modal** joint element the force element is connected to.

If the optional keyword **accelerations** is present, the output will also contain the modal accelerations.

If the optional keyword **type** is present, the alternative keywords **rigid**, **modal** or **all** inform MBDyn about whether only the rigid or the modal portions of the force are expected. The default is **all**, but some communicators might require a specific value. If **rigid** or **all** are used, the **modal** joint element must be connected to a **modal** node. The number of modes of the modal force must match that of the active modes in the **modal** joint element.

File communicator

If **type** is **rigid** or **all**, the first line of the input file must contain:

- the label of the **modal** node the **modal** joint element is connected to, or 0 if the element is grounded;
- three components of force in the **modal** node reference frame;
- three components of moment in the **modal** node reference frame.

If the **modal** joint element is grounded, the force and moment are ignored. If **type** is **modal** or **all**, each of the following lines contains the value of the generalized force for the respective mode.

The output file file is written in textual form. If **type** is **rigid** or **all**, the first line contains:

- the label of the **modal** node the **modal** joint element is connected to, or 0 if the element is grounded;
- the position in the global frame;
- the orientation matrix in the global frame;
- the velocity with respect to the global frame;
- the angular velocity with respect to the global frame.

If the optional keyword **accelerations** is present, at the end of each line the following optional fields will appear:

- the acceleration with respect to the global frame;
- the angular acceleration with respect to the global frame.

If the **modal** joint element is grounded, all the data corresponds to that of the origin of the global reference frame. If **type** is **modal** or **all**, each of the following lines contains, for each mode:

- the value of the generalized variable;
- the value of the first derivative of the generalized variable.

If the optional keyword **accelerations** is present, at the end of each line the following optional field will appear:

- the value of the second derivative of the generalized variable.

EDGE communicator

When the **edge** communicator is used, the **accelerations** keyword is ignored. **type** must be either **rigid** or **modal**. If both need to be used, two **external modal** forces must be instantiated: one with **type** set to **rigid** and the other with **type** set to **modal**.

Rigid. When **type** is **rigid**, the data file written by the external software is expected to contain:

```
* this is a comment
body_forces,R,1,6,0
<fx> <fy> <fz> <mx> <my> <mz>
```

where **fx**, **fy**, **fz**, **mx**, **my**, **mz** respectively are the force and moment components oriented according to the global reference frame. The moments are referred to the point that represents the origin of the FE model reference frame, or to either of the **origin position** or the **origin node** of the **modal** element, if defined.

The data file written by MBDyn contains:

```
* this is a comment
body_dynamics,N,0,0,3
* Body linear velocity in body axes
VRELV,R,1,3,0
<vx> <vy> <vz>
* Body angular velocity in body axes
VRELM,R,1,3,0
<omegax> <omegay> <omegaz>
* Body reference frame cosines (listed by columns)
OMGMAN,R,3,3,0
<R11> <R21> <R31>
<R12> <R22> <R32>
<R13> <R23> <R33>
```

where:

- **vx**, **vy**, **vz** are the components of the velocity of the rigid body, in body axes;
- **omegax**, **omegay**, **omegaz** are the components of the angular velocity of the rigid body, in body axes;
- **R11**, **R12**, **R13**, **R21**, **R22**, **R23**, **R31**, **R32**, **R33** are the components of the orientation matrix.

Modal. When `type` is `modal`, the external software is expected to write a data file containing:

```
* this is a comment
modal_force_flow,R,<n>,1,0
<f1> ... <fn>
```

where `n` is the number of modes (must match the number of active modes in the `modal` joint element), while `f1` to `fn` are the amplitudes of the modal forces.

The data file written by MBDyn contains:

```
* this is a comment
modal_state,N,0,0,2
modal_coordinate,R,1,<n>,0
<q1> ... <qn>
modal_velocity,R,1,<n>,0
<qp1> ... <qpn>
```

where:

- `n` is again the number of modes;
- `q1` to `qn` are the amplitudes of the modal coordinates;
- `qp1` to `qpn` are the amplitudes of the derivatives of the modal coordinates.

Each field can be separated by an arbitrary amount of whitespace, and by at most one “comma” (“,”). Although specifically developed for ease of interfacing with EDGE, this format is open, and thus can be implemented by other software in order to easily couple with MBDyn.

Socket communicator

TBD

Output

An external modal element writes multiple lines at each time step in the `.frc` file.

If `type` is `ridig` or `all`, the line contains:

- in column 1 the label of the element and that of the corresponding node; the format of this field is `element_label#node_label`;
- in columns 2–4 the three components of the force;
- in columns 5–7 the three components of the moment.

If `type` is `modal` or `all`, each line contains:

- in column 1 the label of the element and the corresponding mode; the format of this field is `element_label.mode_number`;
- in column 2 the amplitude of the corresponding modal force.

8.8.13 External modal mapping

This element allows to communicate with an external software that computes forces applied to a set of structural nodes, and may depend on the values of the corresponding variables. Modal forces are actually computed, based on the value of modal variables. The element thus projects the motion of the set of structural nodes in the space of a set of modes before communicating it to the peer. Similarly, it receives generalized forces that are projected back in the space of the structural nodes before assembly. Communication occurs by means of the communicators illustrated earlier for the `external modal` case.

```
<elem_type> ::= force

<force_type> ::= external modal mapping

<force_arglist> ::= <external_force_communicator> ,
  [ reference node , <ref_node_label> , ]
  [ use rigid body forces , { yes | no } ,
    [ rotate rigid body forces , { yes | no } , ] ]
  [ accelerations , { yes | no } , ]
  [ type , { rigid | modal | all } , ]
  nodes number , <num_nodes> ,
  <node_1_label> [ , ... ] ,
  modes number , { from file | <num_modes> } ,
  { full | sparse } mapping file , " <mapping_file_name> "
  [ , threshold , <threshold> ]
```

`ref_node_label` is the label of the reference node if the modes express the motion relative to a specific node.

The mapping file contains the matrix that maps the nodal and modal displacements, velocities and forces. When `full mapping file` is used, a file containing $\text{num_modes} \times (6 \times \text{num_nodes})$ coefficients is expected. When `sparse mapping file` is used, a file containing

```
<row> <col> <value>
```

triplets is expected, with $1 \leq \text{row} \leq \text{num_modes}$ and $1 \leq \text{col} \leq 6 \times \text{num_nodes}$. Both files may start with an arbitrary number of lines beginning with a hash mark (`#`). They are treated as comments and ignored. If a `threshold` is given, only elements whose absolute value is larger than threshold are retained. The value of `threshold` defaults to 0. The mapping matrix is internally stored and handled as sparse, regardless of the file format. When the keyword `from file` is used, the number of modes `num_modes` is computed from the matrix contained in the file `mapping_file_name`.

If the optional keyword `accelerations` is present, the output will also contain the modal accelerations.

If the optional keyword `type` is present, the alternative keywords `rigid`, `modal` or `all` inform MBDyn about whether only the rigid or the modal portions of the force are expected. The default is `all`, but some communicators might require a specific value. If `rigid` or `all` are used, the `reference node` must be defined.

File communicator

Same as `external modal`.

EDGE communicator

Same as `external modal`.

Socket communicator

Same as `external modal`.

Output

An external modal mapping element writes multiple lines at each time step in the `.frc` file.

If `type` is `ridig` or `all`, the line contains:

- in column 1 the label of the element and that of the corresponding node; the format of this field is `element_label@node_label`;
- in columns 2–4 the three components of the force, as received from the peer;
- in columns 5–7 the three components of the moment, as received from the peer;
- in columns 8–10 the three components of the force, in the global reference frame;
- in columns 11–13 the three components of the moment with respect to the reference node, in the global reference frame;
- in columns 14–16 the three components of the force, in the global reference frame, resulting from the combination of the mapped forces;
- in columns 17–19 the three components of the moment with respect to the reference node, in the global reference frame, resulting from the combination of the mapped forces and moments.

If `type` is `modal` or `all`, two types of output are generated:

- a line for each node participating in the mapping, containing:
 - in column 1 the label of the element and the corresponding node; the format of this field is `element_label#node_label`;
 - in columns 2–4 the three components of the corresponding nodal force, in the global reference frame;
 - in columns 5–7 the three components of the corresponding nodal moment, in the global reference frame;
 - in columns 8–10 the three components of the nodal displacement, either in the global reference frame, or in that of the reference node, if any;
 - in columns 11–13 the three components of the nodal orientation, either in the global reference frame, or in that of the reference node, if any;
 - in columns 14–16 the three components of the nodal velocity, either in the global reference frame, or in that of the reference node, if any;
 - in columns 17–19 the three components of the nodal angular velocity, either in the global reference frame, or in that of the reference node, if any.
- a line for each mode the motion is mapped to, containing:
 - in column 1 the label of the element and the corresponding mode; the format of this field is `element_label.mode_number`;
 - in column 2 the amplitude of the corresponding modal force;
 - in column 3 the amplitude of the corresponding modal variable;
 - in column 4 the amplitude of the corresponding modal variable derivative.

8.9 Genel Element

GENEL is the compact form for GENERAL ELEMENT. Those elements that cannot in general be classified in a precise way, or are just under development and thus are not collected in a class of their own until their configuration is stabilized, usually are classified as GENEL. The syntax of the GENEL elements is:

```
<elem_type> ::= genel  
  
<normal_arglist> ::= <genel_type> , <genel_arglist>
```

The output goes in a file with extension `.gen`; only few elements actually generate output. At present, the GENEL class contains very basic general elements and some elements specifically developed for rotorcraft analysis. The latter could be moved to a more specific class in future releases.

8.9.1 General Purpose Elements

Clamp

```
<genel_type> ::= clamp  
  
<genel_arglist> ::=  
  (NodeDof) <clamped_node> ,  
  (DriveCaller) <imposed_value>
```

This element simply forces one arbitrary degree of freedom to assume a value depending on the drive.

Output The format is:

- the label of the element
- the value of the reaction unknown

Distance

```
<genel_type> ::= distance  
  
<genel_arglist> ::=  
  (NodeDof) <node_1> ,  
  (NodeDof) <node_2> ,  
  (DriveCaller) <imposed_distance>
```

This element forces the difference between two arbitrary degrees of freedom to assume the value dictated by the driver.

Output The format is:

- the label of the element
- the value of the reaction unknown

Spring

```
<genel_type> ::= spring

<genel_arglist> ::=
  (NodeDof) <node_1> ,
  (NodeDof) <node_2> ,
  (ConstitutiveLaw<1D>) <const_law>
```

The constitutive law must be **elastic**, but the **distance** genel can apply to arbitrary order degrees of freedom, even between degrees of freedom of different order.

Spring support

```
<genel_type> ::= spring support

<genel_arglist> ::=
  (NodeDof) <node> ,
  (ConstitutiveLaw<1D>) <const_law>
```

The **spring support** must use the **algebraic** value of a **differential** node, but it can use an arbitrary constitutive law, i.e. an elastic constitutive law for a spring, or a viscous constitutive law for a damper, and so on.

Cross spring support

```
<genel_type> ::= cross spring support

<genel_arglist> ::=
  (NodeDof) <row_node> ,
  (NodeDof) <col_node> ,
  (ConstitutiveLaw<1D>) <const_law>
```

It writes a term depending on the **col_node** degree of freedom in an arbitrary manner (given by the **const_law**) to the **row_node** equation.

The **cross spring support** must use the **algebraic** value of a **differential** node, but can use an arbitrary constitutive law, i.e. an elastic constitutive law for a spring, or a viscous constitutive law for a damper, and so on.

Mass

```
<genel_type> ::= mass

<genel_arglist> ::=
  (NodeDof) <node> ,
  (DriveCaller) <mass>
```

The mass must use the **algebraic** value of a **differential** node. The derivative of the **differential** value of the dof is differentiated in a state-space sense, and an inertial driven term is applied to the equation related to the dof:

$$\begin{aligned} m\dot{u} + \dots &= f \\ u - \dot{x} &= 0 \end{aligned}$$

Scalar filter

```

<genel_type> ::= scalar filter

<genel_arglist> ::=
  (NodeDof) <output_node> ,
  { [ node dof , ] (NodeDof) <input_node>
    | drive , (DriveCaller) <input_value> } ,
  [ canonical form , { controllable | observable } , ]
  <output_order> [ , <output_coef_list> ] ,
  <input_order> [ , <input_coef_list> ]
  [ , gain , <gain> ]
  [ , balance , { yes | no } ]

```

This element models a scalar filter of the form

$$A(s)y = B(s)u$$

where A , B are polynomials of arbitrary order, namely

$$y(s) = \frac{b_0 s^{n_n} + b_1 s^{n_n-1} + \dots + b_{n_n}}{a_0 s^{n_d} + a_1 s^{n_d-1} + \dots + a_{n_d}} u(s)$$

The filter must be proper, namely the `output_order` n_d must be greater than or at most equal to the `input_order`, n_n . The polynomial A is assumed to be monic, so the coefficient $a_0 = 1$ and only the coefficients from a_1 to a_{n_d} must be input, while all the coefficients of polynomial B are required, i.e. from b_0 to b_{n_n} .

If a `gain` is supplied, all the coefficients of B are multiplied by the `gain`.

Originally, the input needed to be taken from a `ScalarDof`; now, it can also be taken from a `drive`, which can completely replace the original form (e.g. by means of the `dof`, the `node` and the `element` drives) and requires much less effort because no auxiliary node needs to be set up.

The filter is internally rewritten as a `state space SISO` element. By default, the `controllable` canonical form is used, namely

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \vdots \\ \dot{x}_{n_d} \end{pmatrix} = \begin{bmatrix} -a_1 & -a_2 & -a_3 & \dots & -a_{n_d} \\ 1 & 0 & 0 & & 0 \\ 0 & 1 & 0 & & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_{n_d} \end{pmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} u(t)$$

$$y = [b_1 \quad b_2 \quad b_3 \quad \dots \quad b_{n_d}] \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_{n_d} \end{pmatrix}$$

Otherwise, a **observable** canonical form can be used, namely

$$\begin{aligned} \begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \vdots \\ \dot{x}_{n_d} \end{pmatrix} &= \begin{bmatrix} -a_1 & 1 & 0 & \dots & 0 \\ -a_2 & 0 & 1 & & 0 \\ -a_3 & 0 & 0 & & 0 \\ \vdots & & & \ddots & \vdots \\ -a_{n_d} & 0 & 0 & \dots & 0 \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_{n_d} \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ \vdots \\ b_{n_d} \end{pmatrix} u(t) \\ y &= [1 \ 0 \ 0 \ \dots \ 0] \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_{n_d} \end{pmatrix} \end{aligned}$$

In case the filter is not strictly proper, because $n_n = n_d = n$, it is turned into a strictly proper one by the transformation

$$\begin{aligned} H(s) &= \frac{B(s)}{A(s)} \\ &= \frac{b_0 s^n + b_1 s^{n-1} + \dots + b_n}{s^n + a_1 s^{n-1} + \dots + a_n} \\ &= b_0 + \frac{B(s) - b_0 A(s)}{A(s)} \\ &= b_0 + \frac{(b_1 - b_0 a_1) s^{n-1} + \dots + (b_n - b_0 a_n)}{s^n + a_1 s^{n-1} + \dots + a_n} \end{aligned}$$

and the state space realization is modified by using the resulting numerator coefficients and by introducing a direct feedthrough term $b_0 u(t)$ in the output.

Example. To model a 2nd order Butterworth filter of equation

$$H(s) = \frac{1}{\left(\frac{s}{\omega_c}\right)^2 + \sqrt{2} \left(\frac{s}{\omega_c}\right) + 1} \quad (8.69)$$

use

```
set: real OMEGA_C = 1.*2*pi;           # filter at 1Hz
genel: 1, scalar filter,
      100, abstract, algebraic,        # output node
      drive, sine, 0., 2.*pi, 1., forever, 0., # input drive
      2, sqrt(2.)*OMEGA_C, OMEGA_C^2,   # a_1, a_2
      0, 1.,                             # b_0
      gain, OMEGA_C^2;
```

Note that the formula of Eq. (8.69) had to be rearranged in order to make the denominator monic.

State space SISO

```
<genel_type> ::= state space SISO
```

```

<genel_arglist> ::=
  (NodeDof) <output_node> ,
  { [ node dof , ] (NodeDof) <input_node>
    | drive , (DriveCaller) <input_value> } ,
  <state_order>
  [ , matrix E , <matrix_E_coef_list> ] ,
  matrix A , <matrix_A_coef_list> ,
  matrix B , <matrix_B_coef_list> ,
  matrix C , <matrix_C_coef_list>
  [ , matrix D , <D_coef> ]
  [ , gain , <gain> ]
  [ , balance , { yes | no } ]
  [ , value , <x0_list>
    [ , derivative , <x0p_list> ] ]

```

This element models a scalar (SISO) state space filter of the form

$$\begin{aligned} \mathbf{E}\dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}u \\ y &= \mathbf{C}\mathbf{x} + Du \end{aligned}$$

where

- \mathbf{E} is a(n optional) matrix $\text{state_order} \times \text{state_order}$, defaulting to the identity matrix \mathbf{I} ,
- \mathbf{A} is a matrix $\text{state_order} \times \text{state_order}$,
- \mathbf{B} is a vector $\text{state_order} \times 1$,
- \mathbf{C} is a vector $1 \times \text{state_order}$, and
- D is a(n optional) scalar.

The matrices are read row-oriented; e.g., for matrix \mathbf{A} :

```

matrix A,
  a11, a12, ..., a1N,
  a21, a22, ..., a2N,
  ...,
  aN1, aN2, ..., aNN

```

If LAPACK is available, the matrices are balanced by default, unless explicitly disabled by using the `balance` optional keyword. If matrix \mathbf{E} is defined, the `dggbal()` routine is used; the `dgebal()` routine is used otherwise.

The state and its derivative can be initialized using the keywords `value` and `derivative`. The initialization of the derivative only makes sense when the filter is in descriptor form.

State space MIMO

```

<genel_type> ::= state space MIMO

<genel_arglist> ::=
  <num_outputs> , (NodeDof) <output_node_list> ,
  <num_inputs> ,

```

```

    { [ node dof , ] (NodeDof) <input_node>
      | drive , (DriveCaller) <input_value> }
    [ , ... ] ,
    <state_order>
    [ , matrix E , <matrix_E_coef_list> ] ,
    matrix A , <matrix_A_coef_list> ,
    matrix B , <matrix_B_coef_list> ,
    matrix C , <matrix_C_coef_list>
    [ , matrix D , <matrix_D_coef_list> ]
    [ , gain , <gain> ]
    [ , balance , { yes | no } ]
    [ , value , <x0_list>
      [ , derivative , <x0p_list> ] ]

```

This element models a vector (MIMO) state space filter of the form

$$\begin{aligned} E\dot{x} &= Ax + Bu \\ y &= Cx + Du \end{aligned}$$

where

- E is a(n optional) matrix `state_order` × `state_order`, defaulting to the identity matrix I ,
- A is a matrix `state_order` × `state_order`,
- B is a matrix `state_order` × `num_inputs`,
- C is a matrix `num_outputs` × `state_order`, and
- D is a(n optional) matrix `num_outputs` × `num_inputs`.

The matrices are read row-oriented.

If LAPACK is available, the matrices are balanced by default, unless explicitly disabled by using the `balance` optional keyword. If matrix E is defined, the `dggbal()` routine is used; the `dgebal()` routine is used otherwise.

The state and its derivative can be initialized using the keywords `value` and `derivative`. The initialization of the derivative only makes sense when the filter is in descriptor form.

8.9.2 Special Rotorcraft Genel Elements

The GENEL elements specifically developed for use in rotorcraft analysis currently include:

- the `swashplate`;
- the `rotor trim`.

Swashplate

The `swashplate` GENEL is used to transform the controls of a rotor, in terms of collective and fore/aft and lateral cyclic pitch, into the elongations of the actuators that actually move the swash plate. The syntax of the `swashplate` is:

```

<genel_type> ::= swash plate

<genel_arglist> ::=
  <collective_abstract_node>
    [ , limits , <min_collective> , <max_collective> ] ,
  <fore/aft_abstract_node>
    [ , limits , <min_fore/aft> , <max_fore/aft> ] ,
  <lateral_abstract_node>
    [ , limits , <min_lateral> , <max_lateral> ] ,
  <actuator_1_abstract_node_label> ,
  <actuator_2_abstract_node_label> ,
  <actuator_3_abstract_node_label>
  [ , <dynamic_coef> , <cyclic_factor> , <collective_factor> ]

```

The first three abstract nodes contain the input values of collective, fore/aft and cyclic pitch (they can be actuated by means of abstract forces), and the limits on the “angles” can be easily set. The last three nodes will contain the values of the stroke of the actuators.

The first actuator should be put at the azimuthal position that makes the current blade assume the fore/aft pitch, so that its pitch link is $\text{atan}(1/\omega_\beta)$ before the front direction. The other actuators should follow, 120 deg apart in clockwise direction.

The limits on the actuators will simply force the value of the control inputs to remain in the boundaries regardless of the input values.

The last three optional parameters are a dynamic coefficient that is used to add some dynamics to the actuators’ stroke, namely the input variables are applied a sort of **spring stiffness bulk** element, while the actuators’ strokes are applied a transfer function of the first order, namely $\alpha\dot{x} + x = f$, where $\alpha = \text{dynamic_coef}$ and f is the desired stroke, so the smaller is α , the more the behavior is static.

The **cyclic_factor** and the **collective_factor** parameters are used to scale the inputs from angles in the desired units to strokes, that usually are dimensional parameters. The actual strokes are made of the collective contribution multiplied by **collective_factor**, and the cyclic contribution multiplied by **collective_factor** \times **cyclic_factor**.

Rotor Trim

The syntax of the **rotor trim** is

```

<genel_type> ::= rotor trim

<genel_arglist> ::= [ <rotor_trim_type> ] ,
  <type_specific_rotor_trim_data> ,
  <thrust_node_label> ,
  <longitudinal_moment_node_label> ,
  <lateral_moment_node_label> ,
  (DriveCaller) <desired_thrust_coefficient> ,
  (DriveCaller) <desired_longitudinal_moment_coefficient> ,
  (DriveCaller) <desired_lateral_moment_coefficient> ,
  <rotor_lock_number> ,
  <rotor_nondimensional_flap_frequency> ,
  <thrust_time_constant> , <moments_time_constant> ,
  <thrust_gain> , <moments_gain>
  [ , trigger , (DriveCaller) <trigger> ]

```

```

<rotor_trim_type> ::= { rotor | generic }

# for rotor
<type_specific_rotor_trim_data> ::=
  <rotor_label>

# for generic
<type_specific_rotor_trim_data> ::=
  [ reference node , <ref_node_label> , ]
  (DriveCaller) <thrust> ,
  (DriveCaller) <longitudinal_moment> ,
  (DriveCaller) <lateral_moment> ,
  <rotor_radius> ,
  (DriveCaller) <angular_velocity> ,
  (DriveCaller) <advance_ratio>

```

The **rotor trim** is experimental; it allows to set the controls of a generic helicopter, in conjunction with the **swashplate** element, to asymptotically obtain the desired level of thrust and moment coefficients. The corresponding behavior in terms of trim values (flapping angles and shaft angle) has not been implemented yet. For details, see [15].

8.10 Gravity Element

```

<elem_type> ::= gravity

<arglist> ::= (TplDriveCaller<Vec3>) <gravity_acceleration>

```

the drive of the gravity acceleration, in the global reference frame.

8.11 Hydraulic Element

Note: under development; syntax subjected to changes

Initially implemented by: Lamberto Puggelli

Reviewed by: Pierangelo Masarati

```

<elem_type> ::= hydraulic

<normal_arglist> ::= <hydr_elem_type> , <hydr_elem_data>

```

The field **hydraulic_element_data** usually contains information about fluid properties, which are handled by means of a **HydraulicFluid**. This can be directly inserted, following the syntax described in Section 2.3.4 preceded by the keyword **fluid**, or a previously defined fluid can be recalled by using the keyword **reference** followed by the label of the desired fluid:

```

<hydraulic_fluid_properties> ::=
  { <hydraulic_fluid_specification>
    | reference , <hydraulic_fluid_label> }

```

8.11.1 Actuator

The `actuator` element models the hydraulic and interactional aspects of a two-way linear hydraulic actuator. The two hydraulic nodes represent the pressure in the two chambers of the cylinder. The two structural nodes represent the cylinder and the piston. Their relative motion is assumed to consist essentially in a displacement along a given direction in the reference frame of structural node 1. The user must take care of this by constraining the two nodes with an appropriate set of `joint` elements. Other hydraulic aspects of this component, like leakages between the two chambers and between the chambers and the outside, and the related pressure losses, must be explicitly taken care of, e.g. by means of `minor loss` elements.

```
<hydr_elem_type> ::= actuator

<hydr_elem_data> ::=
  <node_1> , <node_2> ,
  <struct_node_1> , (Vec3) <offset_1> ,
  <struct_node_2> , (Vec3) <offset_2> ,
  [ direction , ((Unit)Vec3) <direction> , ]
  <area_1> ,
  <area_2> ,
  <cylinder_length> ,
  fluid , (HydraulicFluid) <fluid_1> ,
  { same | fluid , (HydraulicFluid) <fluid_2> }
```

The vector `direction` is internally normalized to unity. By default, it is direction 3 in the reference frame of the `structural` node `struct_node_1`.

Example. See the “actuator” example at

<http://www.aero.polimi.it/mbdyn/documentation/examples/actuator>

and the related chapter of the tutorials

<http://www.aero.polimi.it/mbdyn/documentation/tutorials/>

8.11.2 Minor Loss

A pressure loss between two pressure nodes.

```
<hydr_elem_type> ::= minor loss

<hydr_elem_data> ::=
  <node_1> , <node_2> ,
  <k12> , <k21> , <area> ,
  fluid , (HydraulicFluid) <fluid>
```

Coefficients `k12` and `k21` characterize the pressure loss when the flow goes from `node_1` to `node_2` and viceversa. Turbulent flow is assumed.

8.11.3 Three Way Minor Loss

A pressure loss between three pressure nodes, depending on the sign of the pressure drop.

```
<hydr_elem_type> ::= three way minor loss
```

```
<hydr_elem_data> ::=  
  <node_1> , <node_2> , <node_3> ,  
  <k12> , <k31> , <area_12> , <area_31> ,  
  fluid , (HydraulicFluid) <fluid>
```

Coefficients `k12` and `k31` and the respective values of area characterize the pressure loss when the flow goes from `node_1` to `node_2` and from `node_3` to `node_1`, respectively. Turbulent flow is assumed.

8.11.4 Control Valve

```
<hydr_elem_type> ::= control valve
```

```
<hydr_elem_data> ::=  
  <node_1> , <node_2> ,  
  <node_3> , <node_4> ,  
  <area> ,  
  [ loss , <loss_factor> , ]  
  (DriveCaller) <state> ,  
  fluid , (HydraulicFluid) <fluid>
```

This element represents a valve that connects `node_1` to `node_2` and `node_3` to `node_4` when `state` is positive, and `node_1` to `node_3` and `node_2` to `node_4` when `state` is negative. The flow area is proportional to `area` times the norm of `state`, which must be comprised between -1 and 1 . If `loss_factor` is defined, it represents the fraction of area that leaks when `state` is exactly zero.

8.11.5 Dynamic Control Valve

```
<hydr_elem_type> ::= dynamic control valve
```

```
<hydr_elem_data> ::=  
  <node_1> , <node_2> ,  
  <node_3> , <node_4> ,  
  (DriveCaller) <force> ,  
  <initial_displacement> ,  
  <max_displacement> ,  
  <duct_width> ,  
  [ loss , <loss_factor> , ]  
  <valve_diameter> ,  
  <valve_density> ,  
  <displacement_penalty> ,  
  <velocity_penalty> ,  
  <acceleration_penalty> ,  
  fluid , (HydraulicFluid) <fluid>
```

This element represents a valve that connects `node_1` to `node_2` and `node_3` to `node_4` when the displacement is positive and `node_1` to `node_3` and `node_2` to `node_4` when the displacement is negative, accounting for the dynamics of the valve body. The control force `force` is applied to the valve, whose geometric and structural properties are described by `initial_displacement`, `max_displacement`, `duct_width`, `valve_diameter` and `valve_density`. Again the `loss_factor`, if defined, represents the fraction of the

area that leaks when the displacement is zero. Finally, `displacement_penalty`, `velocity_penalty` and `acceleration_penalty` are the penalty coefficients for displacement, velocity and acceleration when the maximum stroke is reached.

8.11.6 Pressure Flow Control Valve

```
<hydr_elem_type> ::= pressure flow control valve

<hydr_elem_data> ::=
  <node_1> , <node_2> ,
  <node_3> , <node_4> ,
  <node_5> , <node_6> ,
  (DriveCaller) <force> ,
  <initial_displacement> ,
  <max_displacement> ,
  <duct_width> ,
  [ loss , <loss_factor> , ]
  <valve_diameter> ,
  <valve_density> ,
  <displacement_penalty> ,
  <velocity_penalty> ,
  <acceleration_penalty> ,
  fluid , (HydraulicFluid) <fluid>
```

Same as Dynamic Control Valve (8.11.5), only the pressures at `node_5` and `node_6` are applied at the sides of the valve body and participate in the force balance.

8.11.7 Pressure Valve

```
<hydr_elem_type> ::= pressure valve

<hydr_elem_data> ::=
  <node_1> , <node_2> ,
  <area> ,
  <mass> ,
  <max_area> ,
  <spring_stiffness> ,
  <spring_preload> ,
  <width> ,
  <displacement_penalty> ,
  <velocity_penalty> ,
  <acceleration_penalty> ,
  fluid , (HydraulicFluid) <fluid>
```

8.11.8 Flow Valve

```
<hydr_elem_type> ::= flow valve

<hydr_elem_data> ::=
  <node_1> , <node_2> , <node_3> ,
```

```

<area> ,
<mass> ,
<max_area> ,
<spring_stiffness> ,
<spring_preload> ,
<width> ,
<displacement_penalty> ,
<velocity_penalty> ,
<acceleration_penalty> ,
fluid , (HydraulicFluid) <fluid>

```

8.11.9 Orifice

```

<hydr_elem_type> ::= orifice

<hydr_elem_data> ::=
  <node_1> , <node_2> ,
  <orifice_area> ,
  [ area , <pipe_area> , ]
  [ reynolds , <critical_reynolds_number> , ]
  fluid , (HydraulicFluid) <fluid>

```

8.11.10 Accumulator

Not documented yet.

8.11.11 Tank

Not documented yet.

8.11.12 Pipe

This element models a simple pipeline connecting two **hydraulic** nodes. In detail, it models the pressure loss due to fluid viscosity according to a constitutive law that depends on the Reynolds number computed on average fluid velocity and hydraulic diameter. Transition between laminar and turbulent flow is also modeled.

```

<hydr_elem_type> ::= pipe

<hydr_elem_data> ::=
  <node_1> , <node_2> ,
  <hydraulic_diameter> ,
  [ area , <area> , ]
  <length> ,
  [ turbulent , ]
  [ initial value , <flow> , ]
  fluid , (HydraulicFluid) <fluid>

```

When **area** is not given, it defaults to the equivalent area computed according to the hydraulic diameter. The flag **turbulent** forces the flow to be considered turbulent since the first iteration, when the initial conditions would fall in the transition region. The **initial value** parameter refers to the initial value of the flow internal state.

Example.

```
set: integer NODE_1 = 10;
set: integer NODE_2 = 20;
set: integer PIPE = 100;
set: integer FLUID = 1000;
# ...
hydraulic: PIPE, pipe, NODE_1, NODE_2,
           5e-3, 100.e-3,
           fluid, reference, FLUID;
```

8.11.13 Dynamic Pipe

Same syntax as the `pipe` hydraulic element (Section 8.11.12), it also considers fluid compressibility in terms of pressure time derivative, and thus the corresponding dynamic effect.

```
<hydr_elem_type> ::= dynamic pipe
```

8.11.14 Imposed Pressure

No specific element has been implemented to impose the pressure at one node. The `clamp` variant of the `genel` element (Section 8.9.1) can be used instead.

Example.

```
set: integer HYDRAULIC_NODE_LABEL = 100;
set: integer IMPOSED_PRESSURE_LABEL = 200;
genel: IMPOSED_PRESSURE_LABEL, clamp,
       HYDRAULIC_NODE_LABEL, hydraulic,
       const, 101325.0;
```

8.11.15 Imposed Flow

No specific element has been implemented to impose the flow at one node. The `abstract` variant of the `force` element (Section 8.8.2) can be used instead. The magnitude of the abstract force is the mass flow extracted from the circuit at that node. In fact, a negative value of the abstract force means that the flow enters the node.

Example.

```
set: integer HYDRAULIC_NODE_LABEL = 100;
set: integer IMPOSED_FLOW_LABEL = 200;
force: IMPOSED_PRESSURE_LABEL, abstract,
       HYDRAULIC_NODE_LABEL, hydraulic,
       const, -1e-3;
```

8.12 Joint Element

Joint elements connect structural nodes. Many different joints are available. Joints may have internal degrees of freedom (the reaction forces) when they introduce kinematic constraints in form of algebraic

relationships between the coordinates of the nodes they connect. Joints that introduce configuration-dependent forces are flexible joints. From the point of view of the input syntax there is no difference between the two classes. A typical joint entry is

```
<elem_type> ::= joint
```

```
<normal_arglist> ::= <joint_type> , <joint_arglist>
```

The output is written to a file with extension `.jnt`. The output is generally made of a standard part, plus some extra information depending on the type of joint, which, when available, is described along with the joint description. Here the standard part is described:

- column 1: the label of the joint
- columns 2–4: the three components of the reaction force in a local reference
- columns 5–7: the three components of the reaction couple in a local frame
- columns 8–10: the three components of the reaction force in the global frame
- columns 11–13: the three components of the reaction couple, rotated into the global frame

Legal joint types, with relative data, are:

8.12.1 Angular acceleration

This joint imposes the absolute angular acceleration of a node about a given axis.

```
<joint_type> ::= angular acceleration
```

```
<joint_arglist> ::=
  <node_label> ,
  ((Unit)Vec3) <relative_direction> ,
  (DriveCaller) <acceleration>
```

The axis is `relative_direction`; it is internally normalized to unity.

Private Data

The following data is available:

1. `"M"` constraint reaction moment along joint direction
2. `"wp"` imposed angular acceleration along joint direction

8.12.2 Angular velocity

This joint imposes the absolute angular velocity of a node about a given axis.

```
<joint_type> ::= angular velocity
```

```
<joint_arglist> ::=
  <node_label> ,
  ((Unit)Vec3) <relative_direction> ,
  (DriveCaller) <velocity>
```

The rotation axis of the imposed angular velocity is `relative_direction`; it is internally normalized to unity. It rotates with the node indicated by `node_label`.

Private Data

The following data is available:

1. "w" imposed angular velocity about joint direction

8.12.3 Axial rotation

This joint is equivalent to a **revolute hinge** (see Section 8.12.38), but the angular velocity about axis 3 is imposed by means of the driver.

```
<joint_type> ::= axial rotation

<joint_arglist> ::=
  <node_1> ,
  [ position , ] (Vec3) <relative_offset_1> ,
  [ orientation , (OrientationMatrix) <relative_orientation_matrix_1> , ]
  <node_2> ,
  [ position , ] (Vec3) <relative_offset_2> ,
  [ orientation , (OrientationMatrix) <relative_orientation_matrix_2> , ]
  (DriveCaller) <angular_velocity>
```

This joint forces `node_1` and `node_2` to rotate about relative axis 3 with imposed angular velocity `angular_velocity`.

Private Data

The following data is available:

1. "rz" relative rotation angle about revolute axis
2. "wz" relative angular velocity about revolute axis
3. "Fx" constraint reaction force in node 1 local direction 1
4. "Fy" constraint reaction force in node 1 local direction 2
5. "Fz" constraint reaction force in node 1 local direction 3
6. "Mx" constraint reaction moment about node 1 local direction 1
7. "My" constraint reaction moment about node 1 local direction 2
8. "Mz" constraint reaction moment about node 1 local direction 3

Hints

When wrapped by a **driven** element, the following hints are honored:

- `hinge{1}` the relative orientation of the joint with respect to node 1 is reset;
- `hinge{2}` the relative orientation of the joint with respect to node 2 is reset;
- `offset{1}` the offset of the joint with respect to node 1 is reset;
- `offset{2}` the offset of the joint with respect to node 2 is reset;
- unrecognized hints are passed to the friction model, if any.

8.12.4 Beam slider

This joint implements a slider, e.g. it constrains a structural node on a string of three-node beams, as discussed in [16].

```
<joint_type> ::= kinematic

<joint_arglist> ::=
  <slider_node_label> ,
  (Vec3) <relative_offset> ,
  [ hinge , (OrientationMatrix) <relative_orientation> ] ,
  [ type , { spherical | classic | spline } , ]
  <beam_number> ,
  <3_node_beam> ,
  { same | (Vec3) <first_node_offset> } ,
  [ hinge , { same | (OrientationMatrix) <first_node_orientation> } , ] ,
  (Vec3) <mid_node_offset> ,
  [ hinge , (OrientationMatrix) <mid_node_orientation> ] ,
  (Vec3) <end_node_offset> ,
  [ hinge , (OrientationMatrix) <end_node_orientation> ] ,
  [ ... ]
  [ , initial beam , <initial_beam> ]
  [ , initial node , <initial_node> ]
  [ , smearing , <smearing_factor> ]
```

There are three types of slider:

- the **spherical** slider does not constrain the orientation of the node;
- the **classical** slider does allow rotation only about the sliding line;
- the **spline** slider constrain the orientation of the node.

For each node of each beam element, the offset and the orientation of the slider can be defined; except for the first element, the offset and the orientation of the first node can be specified using the keyword **same**, which causes the node to take the same value of the last node of the previous beam. The **initial_beam** and **initial_node** indices serve as hints to set the initial contact point of the sliding node. The **smearing_factor** determines the (non-dimensional) extension of the interference segment when the node passes from one segment to another. Further information is available in [16].

8.12.5 Brake

This element models a wheel brake, i.e. a constraint that applies a frictional internal torque between two nodes about an axis. The frictional torque depends on the normal force that is applied as an external input by means of the same friction models implemented for regular joints.

```
<joint_type> ::= brake

<joint_arglist> ::=
  <node_1_label> , (Vec3) <relative_offset_1>
  [ , hinge , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label> , (Vec3) <relative_offset_2>
  [ , hinge , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
```

```

friction , <average_radius> ,
  [ preload , <const_value> , ]
  <friction_model> ,
  <shape_function> ,
  (DriveCaller) <normal_force>

```

Note: a *revolute hinge* (see Section 8.12.38) between the same two nodes must be defined as well, such that the only allowed relative kinematics between the two nodes are a rotation about relative axis 3.

Output

In addition to the standard output, the **brake** joint outputs:

- the three Euler-Cardano angles that express the relative rotation between the two connected nodes, in degrees;
- the three components of the relative angular velocity, in the reference frame of node 2;
- any output specific for the friction model in use;
- the value of the force the brake is activated with.

Private Data

The following data is available:

1. "rz" relative rotation angle about brake axis
2. "wz" relative angular velocity about brake axis

8.12.6 Cardano hinge

This joint implements a Cardano's joint, also known as Hooke's joint or Universal joint, which is made of a sequence of two revolute hinges orthogonal to each other, one about relative axis 2 and one about relative axis 3 of the reference systems defined by the two **orientation** statements. In other words, this joint constrains the relative axis 3 of node 1 to be always orthogonal to the relative axis 2 of node 2. As a result, torque is transmitted about axis 1 of both nodes. The relative position is constrained as well.

```

<joint_type> ::= cardano hinge

<joint_arglist> ::=
  <node_1_label> ,
  position , (Vec3) <relative_offset_1>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label> ,
  position , (Vec3) <relative_offset_2>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]

```

Note: this joint does not represent a constant velocity joint, so, when a steady deflection between the two nodes is present, a constant velocity about axis 1 of one node results in an oscillating velocity about axis 1 for the other node. The constant velocity joint is implemented by the **gimbal joint** (Section 8.12.22), which essentially consists of a sequence of two **cardano hinge** (see Section 8.12.6) joints with the hinges in reversed order, assembled in a manner that the deflection is split half and half between the two joints.

8.12.7 Cardano pin

This joint implements a “Cardano” joint between a node and the ground. The absolute position is also constrained. See above for details about the formulation.

```
<joint_type> ::= cardano pin

<joint_arglist> ::=
  <node_label> ,
  position , (Vec3) <relative_offset>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix> ] ,
  position , (Vec3) <absolute_pin_position>
  [ , orientation , (OrientationMatrix) <absolute_pin_orientation_matrix> ]
```

*Note: this is equivalent to a **cardano hinge** (see Section 8.12.6) when one node is grounded.*

8.12.8 Cardano rotation

This joint implements a “Cardano” joint, which is made of a sequence of two revolute hinges orthogonal to each other. The relative position is not constrained. See the **cardano hinge** for more details.

```
<joint_type> ::= cardano rotation

<joint_arglist> ::=
  <node_1_label>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]
```

*Note: this is equivalent to a **cardano hinge** (see Section 8.12.6) without position constraint. Or, in other words, a **cardano hinge** (see Section 8.12.6) is made of a **cardano rotation** on top of a **spherical hinge** (see Section 8.12.43).*

8.12.9 Clamp

This joint grounds all 6 degrees of freedom of a node in an arbitrary position and orientation that remains fixed.

```
<joint_type> ::= clamp

<joint_arglist> ::= <node_label> ,
  [ position , ]
  { node | (Vec3) <absolute_position> } ,
  [ orientation , ]
  { node | (OrientationMatrix) <absolute_orientation_matrix> }
```

The keyword **node** forces the joint to use the node’s position and reference frame. Otherwise, they must be entered in the usual way for these entities. The default value for **position** and **orientation** are the position and the orientation of the clamped node.

Private Data

The following data is available:

1. "Fx" constraint reaction force in global direction 1
2. "Fy" constraint reaction force in global direction 2
3. "Fz" constraint reaction force in global direction 3
4. "Mx" constraint reaction moment in local direction 1
5. "My" constraint reaction moment in local direction 2
6. "Mz" constraint reaction moment in local direction 3

8.12.10 Coincidence

not implemented yet, use a `spherical hinge` (see Section 8.12.43) and a `prismatic` instead.

8.12.11 Deformable Axial Joint

This joint implements a configuration dependent moment that is exchanged between two nodes about an axis rigidly attached to the first node. It is intended to be used in conjunction with another joint that constrains the relative rotation between the two nodes about the other axes, although no check is in place.

```
<joint_type> ::= deformable axial joint

<joint_arglist> ::=
  <node_1_label>
    [ , position , (Vec3) <relative_offset_1> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
    [ , position , (Vec3) <relative_offset_2> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
  <(ConstitutiveLaw<1D>> <const_law>
```

Private Data

The following data is available:

1. "rz" relative rotation about axis 3, used as strain in the constitutive law
2. "wz" relative angular velocity about axis 3, used as strain rate in the constitutive law
3. "Mz" moment about axis 3, output by the constitutive law

In addition, the joint provides access to any private data provided by the constitutive law. They are accessed by prefixing the name of the data with the string `"constitutiveLaw."`; see the specific constitutive law description of the available data in Section 2.9.

Hints

When wrapped by a `driven` element, the `deformable axial joint` honors the following hints:

- `hinge{1}` the orientation with respect to node 1 of the reference used to compute the linear strain is reset to the value resulting from the node 2 hinge orientation and the node 1 orientation

$$\tilde{\mathbf{R}}_{1h} = \mathbf{R}_1^T \mathbf{R}_2 \tilde{\mathbf{R}}_{2h}$$

- `hinge{2}` the orientation with respect to node 2 of the reference used to compute the linear strain is reset to the value resulting from the node 1 hinge orientation and the node 2 orientation

$$\tilde{\mathbf{R}}_{2h} = \mathbf{R}_2^T \mathbf{R}_1 \tilde{\mathbf{R}}_{1h}$$

- unrecognized hints are passed through to the constitutive law.

Output

In addition to the standard output, the `deformable axial joint` outputs:

- column 14: the relative rotation about axis 3

If the constitutive law is either viscous or viscoelastic:

- column 15: the relative angular velocity about axis 3

8.12.12 Deformable displacement hinge

Deprecated; use the `deformable displacement joint` (see Section 8.12.13) instead.

8.12.13 Deformable displacement joint

This joint implements a configuration dependent force that is exchanged between two points associated to two nodes with an offset. The force may depend, by way of a generic 3D constitutive law, on the relative position and velocity of the two points, expressed in the reference frame of node 1.

The constitutive law is attached to the reference frame of node 1, so the sequence of the connections may matter in case of anisotropic constitutive laws, if the relative orientation of the two nodes changes during the analysis.

```
<joint_type> ::= deformable displacement joint

<joint_arglist> ::=
  <node_1_label> ,
    position , (Vec3) <relative_offset_1>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label> ,
    position , (Vec3) <relative_offset_2>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
  <(ConstitutiveLaw<3D>> <const_law>
```

Note: a variant of this element is under development, which refers the material reference frame to an orientation that is intermediate between those of the two nodes. See the `invariant deformable displacement joint` (Section 8.12.27).

Private Data

The following data is available:

1. **"dx"** relative displacement in node 1 local direction 1
2. **"dy"** relative displacement in node 1 local direction 2
3. **"dz"** relative displacement in node 1 local direction 3
4. **"vx"** relative velocity in node 1 local direction 1
5. **"vy"** relative velocity in node 1 local direction 2
6. **"vz"** relative velocity in node 1 local direction 3
7. **"Fx"** constraint reaction force in node 1 local direction 1
8. **"Fy"** constraint reaction force in node 1 local direction 2
9. **"Fz"** constraint reaction force in node 1 local direction 3

In addition, the joint provides access to any private data provided by the constitutive law. They are accessed by prefixing the name of the data with the string **"constitutiveLaw."**; see the specific constitutive law description of the available data in Section 2.9.

Hints

When wrapped by a **driven** element, the **deformable displacement joint** honors the following hints:

- **hinge{1}** the orientation with respect to node 1 of the reference used to compute the linear strain is reset to the value resulting from the node 2 hinge orientation and the node 1 orientation

$$\tilde{\mathbf{R}}_{1h} = \mathbf{R}_1^T \mathbf{R}_2 \tilde{\mathbf{R}}_{2h}$$

- **hinge{2}** the orientation with respect to node 2 of the reference used to compute the linear strain is reset to the value resulting from the node 1 hinge orientation and the node 2 orientation

$$\tilde{\mathbf{R}}_{2h} = \mathbf{R}_2^T \mathbf{R}_1 \tilde{\mathbf{R}}_{1h}$$

- **offset{1}** the offset of the joint with respect to node 1 is reset to the value resulting from the node 2 offset and the node 1 position

$$\tilde{\mathbf{f}}_1 = \mathbf{R}_1^T \left(\mathbf{x}_2 + \mathbf{R}_2 \tilde{\mathbf{f}}_2 - \mathbf{x}_1 \right)$$

- **offset{2}** the offset of the joint with respect to node 2 is reset to the value resulting from the node 1 offset and the node 2 position

$$\tilde{\mathbf{f}}_2 = \mathbf{R}_2^T \left(\mathbf{x}_1 + \mathbf{R}_1 \tilde{\mathbf{f}}_1 - \mathbf{x}_2 \right)$$

- unrecognized hints are passed through to the constitutive law.

Output

In addition to the standard output, the **deformable displacement joint** outputs:

- columns 14–16: the three components of the linear strain.

If the constitutive law is either viscous or viscoelastic:

- columns 17–19: the three components of the linear strain rate.

8.12.14 Deformable Hinge

This joint implements a configuration dependent moment that is exchanged between two nodes. The moment may depend, by way of a generic 3D constitutive law, on the relative orientation and angular velocity of the two nodes, expressed in the reference frame of node 1.

The constitutive law is attached to the reference frame of node 1, so the sequence of the connections may matter in case of anisotropic constitutive laws, if the relative orientation of the two nodes changes during the analysis.

```
<joint_type> ::= deformable hinge

<joint_arglist> ::=
  <node_1_label>
    [ , position , (Vec3) <relative_position_1> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
    [ , position , (Vec3) <relative_position_2> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
  (ConstitutiveLaw<3D>) <const_law>
  [ , orientation description ,
    { euler123 | euler313 | euler 321
      | orientation vector | orientation matrix } ]
```

Note: a variant of this element has been developed, which refers the material reference frame to an orientation that is intermediate between those of the two nodes. See the **invariant deformable hinge** (Section 8.12.28).

Note: this joint only applies internal moments; no forces and no constraints to displacements are applied. Usually, it is best used in conjunction with a **spherical hinge** (see Section 8.12.43), or any kind of joint that constrains the relative displacement of the nodes as appropriate.

Private Data

The following data is available:

1. **"rx"** relative rotation vector component in node 1 local direction 1
2. **"ry"** relative rotation vector component in node 1 local direction 2
3. **"rz"** relative rotation vector component in node 1 local direction 3
4. **"wx"** relative angular velocity in node 1 local direction 1
5. **"wy"** relative angular velocity in node 1 local direction 2
6. **"wz"** relative angular velocity in node 1 local direction 3

7. "Mx" constraint reaction moment in node 1 local direction 1
8. "My" constraint reaction moment in node 1 local direction 2
9. "Mz" constraint reaction moment in node 1 local direction 3

In addition, the joint provides access to the private data provided by the constitutive law. They are accessed by prefixing the name of the data with the string "constitutiveLaw."; see the specific constitutive law description of the available data in Section 2.9.

Hints

When wrapped by a `driven` element, the `deformable hinge` joint honors the following hints:

- `hinge{1}` the orientation with respect to node 1 of the reference used to compute the angular strain is reset to the value resulting from the node 2 hinge orientation and the node 1 orientation

$$\tilde{\mathbf{R}}_{1h} = \mathbf{R}_1^T \mathbf{R}_2 \tilde{\mathbf{R}}_{2h}$$

- `hinge{2}` the orientation with respect to node 2 of the reference used to compute the angular strain is reset to the value resulting from the node 1 hinge orientation and the node 2 orientation

$$\tilde{\mathbf{R}}_{2h} = \mathbf{R}_2^T \mathbf{R}_1 \tilde{\mathbf{R}}_{1h}$$

- unrecognized hints are passed through to the constitutive law.

Output

In addition to the standard output, the `deformable joint` outputs:

- columns 14–16 (or 14–22): the angular strain in the selected orientation format (either three or nine columns, according to the selected `orientation description`, or to its default value, see Section 5.2.10).

If the constitutive law is either viscous or viscoelastic:

- columns 17–19 (or 23–25): the three components of the angular strain rate.

8.12.15 Deformable joint

This joint implements a configuration dependent force and moment that are exchanged by two points associated to two nodes with an offset. The force and moment may depend, by way of a generic 6D constitutive law, on the relative position, orientation, velocity and angular velocity of the two points, expressed in the reference frame of node 1. It may be thought of as a combination of a `deformable displacement joint` (see Section 8.12.13) and of a `deformable hinge` (see Section 8.12.14); a significant difference is that the `deformable joint` uses a 6D `constitutive law` (see Section 2.9), so the force and the moment may both depend on the displacement and on the orientation.

The constitutive law is attached to the reference frame of node 1, so the sequence of the connections may matter in case of anisotropic constitutive laws, if the relative orientation of the two nodes changes during the analysis.

This element may add a considerable overhead because of the computation of the cross-coupling effects between the forces and moments and the relative positions and orientations; if the constitutive law does not couple relative displacements and orientations, a better choice consists in combining a `deformable hinge` and a `deformable displacement joint`.

```

<joint_type> ::= deformable joint

<joint_arglist> ::=
  <node_1_label> ,
    position , (Vec3) <relative_offset_1>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label> ,
    position , (Vec3) <relative_offset_2>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
  (ConstitutiveLaw<6D>) <const_law>
  [ , orientation description ,
    { euler123 | euler313 | euler 321
      | orientation vector | orientation matrix } ]

```

Private Data

The following data is available:

1. "dx" relative displacement in node 1 local direction 1
2. "dy" relative displacement in node 1 local direction 2
3. "dz" relative displacement in node 1 local direction 3
4. "rx" relative rotation vector component in node 1 local direction 1
5. "ry" relative rotation vector component in node 1 local direction 2
6. "rz" relative rotation vector component in node 1 local direction 3
7. "vx" relative velocity in node 1 local direction 1
8. "vy" relative velocity in node 1 local direction 2
9. "vz" relative velocity in node 1 local direction 3
10. "wx" relative angular velocity in node 1 local direction 1
11. "wy" relative angular velocity in node 1 local direction 2
12. "wz" relative angular velocity in node 1 local direction 3
13. "Fx" constraint reaction force in node 1 local direction 1
14. "Fy" constraint reaction force in node 1 local direction 2
15. "Fz" constraint reaction force in node 1 local direction 3
16. "Mx" constraint reaction moment in node 1 local direction 1
17. "My" constraint reaction moment in node 1 local direction 2
18. "Mz" constraint reaction moment in node 1 local direction 3

In addition, the joint provides access to the private data provided by the constitutive law. They are accessed by prefixing the name of the data with the string "constitutiveLaw."; see the specific constitutive law description of the available data in Section 2.9.

Hints

When wrapped by a **driven** element, the **deformable joint** honors the same hints of the **deformable displacement joint**.

Output

In addition to the standard output, the **deformable joint** outputs:

- columns 14–16: the three components of the linear strain;
- columns 17–19 (or 17–25): the angular strain in the selected orientation format (either three or nine columns, according to the selected **orientation description**, or to its default value, see Section 5.2.10).

If the constitutive law is either viscous or viscoelastic:

- columns 20–22 (or 26–28): the three components of the linear strain rate;
- columns 23–25 (or 29–31): the three components of the angular strain rate.

8.12.16 Distance

This joint forces the distance between two points, each relative to a node, to assume the value indicated by the drive. If no offset is given, the points are coincident with the node themselves.

```
<joint_type> ::= distance

<joint_arglist> ::=
  <node_1_label> ,
  [ position , <relative_offset_1> , ]
  <node_2_label> ,
  [ position , <relative_offset_2> , ]
  { (DriveCaller) <distance> | from nodes }
```

The **relative_offset_*** are the distances of each end of the joint from the relative nodes in the node reference frame. The **distance** and the **distance with offset** joints do not allow null distance.

Note: in case the keyword **from nodes** is used, a constant drive caller is automatically instantiated for the **distance**. Its value is computed from the initial positions of the nodes; if defined, the distance between the offsets is considered.

Output

The extra output is:

- the three components of the imposed distance in the global frame
- the norm of the imposed distance

Private Data

The following data is available:

1. **"d"** enforced distance

Hints

When wrapped by a **driven** element, the following hints are honored:

- unrecognized hints are passed through to the distance drive.

8.12.17 Distance with offset

This element has been deprecated in favor of the **distance** joint element, which now supports offsets. It may be dropped in future releases.

Private Data

See the **distance** joint element.

8.12.18 Drive displacement

This joint imposes the relative position between two points optionally offset from two structural nodes, in the form of a vector that expresses the direction of the displacement in the reference frame of node 1, whose amplitude is defined by a drive.

```
<joint_type> ::= drive displacement

<joint_arglist> ::=
  <node_1_label> , (Vec3) <offset_1> ,
  <node_2_label> , (Vec3) <offset_2> ,
  (TplDriveCaller<Vec3>) <relative_position>
```

This element is superseded by the **total joint**, see Section 8.12.46.

Private Data

The following data is available:

1. **"dx"** imposed displacement component along node 1 axis 1
2. **"dy"** imposed displacement component along node 1 axis 2
3. **"dz"** imposed displacement component along node 1 axis 3
4. **"fx"** reaction force component along node 1 axis 1
5. **"fy"** reaction force component along node 1 axis 2
6. **"fz"** reaction force component along node 1 axis 3

Hints

When wrapped by a **driven** element, the following hints are honored:

- **offset{1}** the offset of the joint with respect to node 1 is reset by pointing to the drive displacement point;
- **offset{2}** the offset of the joint with respect to node 2 is reset by pointing to the drive displacement point;
- unrecognized hints are passed through to the **TplDriveCaller<Vec3>**.

8.12.19 Drive displacement pin

This joint imposes the absolute position of a point optionally offset from a structural node, in the form of a vector that expresses the direction of the displacement in the absolute reference frame, whose amplitude is defined by a drive.

```
<joint_type> ::= drive displacement pin

<joint_arglist> ::=
  <node_label> , (Vec3) <node_offset> ,
  (Vec3) <offset> ,
  (TplDriveCaller<Vec3>) <position>
```

This element is superseded by the `total pin joint`, see Section 8.12.47.

Private Data

The following data is available:

1. `"dx"` imposed displacement component along absolute axis 1
2. `"dy"` imposed displacement component along absolute axis 2
3. `"dz"` imposed displacement component along absolute axis 3
4. `"fx"` reaction force component along absolute axis 1
5. `"fy"` reaction force component along absolute axis 2
6. `"fz"` reaction force component along absolute axis 3

Hints

When wrapped by a `driven` element, the following hints are honored:

- `offset{1}` the offset of the joint with respect to the node is reset by pointing to the driven displacement point;
- `offset{0}` the offset of the joint with respect to the absolute reference frame is reset by pointing to the driven displacement point;
- unrecognized hints are passed through to the `TplDriveCaller<Vec3>`.

8.12.20 Drive hinge

This joint imposes the relative orientation between two nodes, in the form of a rotation about an axis whose amplitude is defined by a drive.

```
<joint_type> ::= drive hinge

<joint_arglist> ::=
  <node_1_label>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
  (TplDriveCaller<Vec3>) <hinge_orientation>
```

This element is superseded by the **total joint**, see Section 8.12.46.

Note: this element is experimental; now it is more reliable, but it is limited to $\| \text{hinge_orientation} \| < \pi$.

Private Data

The following data is available:

1. **"rx"** imposed relative rotation about node 1 hinge axis 1
2. **"ry"** imposed relative rotation about node 1 hinge axis 2
3. **"rz"** imposed relative rotation about node 1 hinge axis 3
4. **"Mx"** reaction moment about node 1 hinge axis 1
5. **"My"** reaction moment about node 1 hinge axis 2
6. **"Mz"** reaction moment about node 1 hinge axis 3

Hints

When wrapped by a **driven** element, the following hints are honored:

- **hinge{1}** the orientation with respect to node 1 of the reference in which the enforced orientation is expressed is reset;
- **hinge{2}** the orientation with respect to node 2 of the reference in which the enforced orientation is expressed is reset;
- unrecognized hints are passed through to the **Vec3** drive.

8.12.21 Gimbal hinge

This joint has not been implemented yet; use a **gimbal rotation** and a **spherical hinge** (see Section 8.12.43) to emulate.

8.12.22 Gimbal rotation

A homokinetic joint without position constraints; this joint, in conjunction with a **spherical hinge** (see Section 8.12.43) joint, should be used to implement an ideal tiltrotor gimbal instead of a **cardano rotation** (see Section 8.12.8). See the technical manual and [17] for details. It is equivalent to a series of two Cardano's joints (the **cardano hinge**, see Section 8.12.6) rotated 90 degrees apart, each accounting for half the relative rotation between axis 3 of each side of the joint.

```
<joint_type> ::= gimbal rotation

<joint_arglist> ::=
  <node_1_label>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]
  [ , orientation description ,
    { euler123 | euler313 | euler321
      | orientation vector | orientation matrix } ]
```

This joint allows nodes 1 and 2 to rotate about relative axes 1 and 2.

Output

The **gimbal rotation** joint outputs the three components of the reaction couple in the local frame (that of node 1) and in the global frame.

The extra columns 14 and 15 contain the angles ϑ and φ . The extra columns from 16 on contain the relative rotation between nodes 1 and 2, in the format determined either by an explicit selection of the orientation description, or by the default orientation description defined by the **default orientation** keyword (see Section 5.2.10).

The three components of the reaction couple in the local frame, and the angles ϑ and φ are also available as private data of the element under the names **lambda[<i>]**, with **i**=1,2,3, **theta** and **phi**. For a description of the formulation and of the angles describe above, see the Technical Manual.

Private Data

The following data is available:

1. **"lambda[1]"** constraint reaction moment about node 1 local axis 1
2. **"lambda[2]"** constraint reaction moment about node 1 local axis 2
3. **"lambda[3]"** constraint reaction moment about node 1 local axis 3
4. **"theta"** relative angle ϑ
5. **"phi"** relative angle φ

8.12.23 Imposed displacement

This joint imposes the relative position between two points, optionally offset from two structural nodes, along a given direction that is rigidly attached to the first node. The amplitude of the displacement is defined by a drive.

```
<joint_type> ::= imposed displacement

<joint_arglist> ::=
  <node_1_label> , (Vec3) <offset_1> ,
  <node_2_label> , (Vec3) <offset_2> ,
  (Vec3) <direction> ,
  (DriveCaller) <relative_position>
```

This element is superseded by the **total joint**, see Section 8.12.46.

Private Data

The following data is available:

1. **"d"** imposed displacement along **direction**, in node 1 reference frame;
2. **"f"** reaction force along **direction**, in node 1 reference frame.

Hints

When wrapped by a `driven` element, the following hints are honored:

- `offset{1}` the offset of the joint with respect to node 1 is reset by pointing exactly to the imposed displacement point;
- `offset{2}` the offset of the joint with respect to node 2 is reset by pointing exactly to the imposed displacement point;
- unrecognized hints are passed through to the drive.

8.12.24 Imposed displacement pin

This joint imposes the absolute displacement of a point optionally offset from a structural node, along a direction defined in the absolute reference frame. The amplitude of the displacement is defined by a drive.

```
<joint_type> ::= imposed displacement pin

<joint_arglist> ::=
  <node_label> , (Vec3) <node_offset> ,
  (Vec3) <offset> ,
  (Vec3) <direction> ,
  (DriveCaller) <position>
```

This element is superseded by the `total pin joint`, see Section 8.12.47.

Private Data

The following data is available:

1. "x" imposed displacement along `direction`, in the absolute reference frame;
2. "f" reaction force along `direction`, in the absolute reference frame.

Hints

When wrapped by a `driven` element, the following hints are honored:

- `offset{1}` the offset of the joint with respect to the node is reset by pointing exactly to the imposed displacement point;
- `offset{0}` the offset of the joint with respect to the absolute reference frame is reset by pointing exactly to the imposed displacement point;
- unrecognized hints are passed through to the drive.

8.12.25 In line

This joint forces a point relative to the second node to move along a line attached to the first node.

```
<joint_type> ::= in line
```

```
<joint_arglist> ::=  
  <node_1_label> ,  
  [ position , (Vec3) <relative_line_position> , ]  
  [ orientation , (OrientationMatrix) <relative_orientation> , ]  
  <node_2_label>  
  [ , offset , (Vec3) <relative_offset> ]
```

A point, optionally offset by `relative_offset` from the position of node `node_2_label`, slides along a line that passes through a point that is rigidly offset by `relative_line_position` from the position of `node_1_label`, and is directed as direction 3 of `relative orientation`.

8.12.26 In plane

This joint forces a point relative to the second node to move in a plane attached to the first node.

```
<joint_type> ::= in plane
```

```
<joint_arglist> ::=  
  <node_1_label> ,  
  [ position , (Vec3) <relative_plane_position> , ]  
  ((Unit)Vec3) <relative_normal_direction> ,  
  <node_2_label>  
  [ , offset , (Vec3) <relative_offset> ]
```

A point, optionally offset by `relative_offset` from the position of node `node_2_label`, slides on a plane that passes through a point that is rigidly offset by `relative_plane_position` from the position of `node_1_label`, and is perpendicular to `relative_normal_direction`. The vector `relative_normal_direction` is internally normalized to unity.

8.12.27 Invariant deformable displacement joint

Under development; right now, use the `deformable displacement joint` (see Section 8.12.13) instead.

8.12.28 Invariant deformable hinge

This (experimental) element is a variant of the `deformable hinge` (see Section 8.12.14); refer to that joint for the input syntax.

The *invariant* form of the `deformable hinge` joint refers the constitutive law to an orientation that is intermediate between those of the nodes it connects. As a result, the moment exchanged between the two nodes is invariant to the sequence of definition of the nodes when anisotropic constitutive laws are used. It is worth stressing that determining the constitutive law for this element may be tricky, since usual measurement approaches directly measure forces and moments in a reference frame that is attached to one of the two ends of the component, so in practical cases it might be more appropriate to use the *variant* form of the joint, and consistently referring the constitutive law to the same end used to measure the mechanical properties of the component.

8.12.29 Kinematic

This joint will eventually evolve into a connection with external programs that impose the entire motion of a node; at present, you can use it by directly writing code in files `mbdyn/struct/kin.cc` and `mbdyn/struct/kin.h` that implement the motion you wish to impose.

```
<joint_type> ::= kinematic

<joint_arglist> ::=
  <node_1_label> ,
  (DriveCaller) <input>
```

This element is superseded by the **total pin joint**, see Section 8.12.47.

Note: this element is eXperimental.

8.12.30 Linear acceleration

This joint imposes the absolute linear acceleration of a node along a given axis.

```
<joint_type> ::= linear acceleration

<joint_arglist> ::=
  <node_label> ,
  ((Unit)Vec3) <relative_direction> ,
  (DriveCaller) <acceleration>
```

The axis is `relative_direction`; it is internally normalized to unity.

Private Data

The following data is available:

1. "F" constraint reaction force along joint direction
2. "a" imposed acceleration along joint direction

8.12.31 Linear velocity

This joint imposes the absolute linear velocity of a node along a given axis.

```
<joint_type> ::= linear velocity

<joint_arglist> ::=
  <node_label> ,
  ((Unit)Vec3) <relative_direction> ,
  (DriveCaller) <velocity>
```

The axis is `relative_direction`; it is internally normalized to unity.

Private Data

The following data is available:

1. "v" imposed velocity along joint direction

8.12.32 Modal

Original implementation: Felice Felippone;

Initial review: Giuseppe Quaranta;

Current review: Pierangelo Masarati.

This joint implements a Component Mode Synthesis (CMS) deformable body. Its interface with the multibody domain is represented by clamps that constrain the multibody interface nodes to the position and orientation of the corresponding FEM nodes.

```
<joint_type> ::= modal

<joint_arglist> ::=
  { <reference_modal_node> | clamped
    [ , position , (Vec3) <absolute_position> ]
    [ , orientation , (OrientationMatrix) <absolute_orientation> ] } ,
  <mode_number> ,
  [ list , <mode> [ , ... ] , ]
  [ initial value ,
    { mode , <n> , <q_n> , <qP_n> [ , ... ]
      | <q_1> , <qP_1> [ ... ] } , ]
  { <fem_node_number> | from file } ,
  [ { no damping
    | proportional damping , <damping_coef>
    | diag damping ,
      { all , <damping_coef> [ , ... ]
        | <damped_modes> , <mode_damping> [ , ... ] } } , ]
  " <fem_data_file> " ,
  [ [ { mass | stiffness } ] threshold , <threshold> , ]
  [ { create binary | use binary | update binary }
    [ , ... ] , ]
  [ use invariant 9 , ]
  [ { origin node , <origin_node> | origin position , (Vec3) <pos> } , ]
  <interface_nodes_number> , <interface_node> [ , ... ]

<mode_damping> ::= <mode_index> , <mode_damping_coef>

<interface_node> ::=
  { " <fem_node_label> " | (Integer) <fem_node_label> } ,
  <multibody_label> , (Vec3) <offset_of_FEM_node>
```

The `reference_modal_node` is a special dynamic structural node that is required to handle the rigid body motion of the modal joint. Its input is completely analogous to that of the `dynamic` structural nodes, see Section 6.1, only the keyword `dynamic` must be replaced by `modal`.

If no rigid body dynamics is required, e.g. if the modal element is clamped, the `clamped` option can be used, which allows to set the optional `absolute_position` and `absolute_orientation`. The former is the location, in the multibody global reference system, of the origin of the FEM reference system. The latter is the orientation, in the multibody global reference frame, of the FEM reference system. They default to zero and identity, which means that the global multibody reference frame and the FEM reference system are coincident.

The mode count in `mode_number` is not required to match the number of modes listed in the FEM data file; if a lower number is given, only the first `mode_number` modes are used; moreover, a list of active

modes can be given, to use non-consecutive modes; e.g., to use modes 1, 2, 3 and 6 out of a set of 10 modes defined in the FEM data file, write

```

... ,
4,                # number of active modes
  list, 1, 2, 3, 6,  # list of active modes
...

```

By default, the origin of the FEM grid is placed either in the position of the modal node, with its orientation, or in the absolute position and orientation in case the modal element is clamped, unless one of the mutually exclusive `origin node` and `origin position` optional keywords is used. The `origin node` keyword defines what FEM node corresponds to the `modal` node in the multibody domain, or what FEM node corresponds to the optional absolute position and orientation in case of a clamped modal element. The `origin position` keyword defines where, in the FEM model reference frame, the `modal` node is placed, or the location that corresponds to the optional absolute position and orientation in case of a clamped modal element.

The FEM labels can be strings made of any character, provided they do not contain blanks. The strings must be enclosed in double quotes. For legacy reasons, a label not enclosed in double quotes is accepted, provided it is an integer.

If any of `mass threshold`, `stiffness threshold` are given, elements whose absolute value is lower than the threshold are discarded from the mass and stiffness matrix, respectively. The keyword `threshold` is a catchall for both, meaning that the same threshold must be used for both matrices.

The list of matchings between FEM and multibody nodes needs special care. The `offset_of_FEM_node` field contains the distance of the FEM node from the respective multibody node; by default, this is expressed in the reference frame of the multibody node.

The `fem_data_file` can be generated by NASTRAN, following the procedure illustrated in Appendix B.3.

It is strongly recommended that constrained modal analysis be used for otherwise free bodies, with the statically determined constraint consisting of clamping the FEM node that will coincide with the node indicated as `reference_modal_node`, and using the `origin node` to make that point the origin of the FEM frame.

Note about reference frames: the coincidence constraint between multibody and FEM nodes is written between the local frame of the FEM node and the global frame of the multibody node. As such, the multibody nodes must be oriented as the `modal` node the `modal` joint refers to, if any, or as the reference orientation of the `modal` joint, if it is clamped.

Note about initial assembly: it is very important that multibody and FEM nodes at interfaces are given with a high degree of accuracy, because the initial assembly procedure of the modal element does not behave very well (it's on the TODO list with a very low priority); as a consequence, pay very much attention to the input of these nodes, until more robust procedures are developed. One trick is to build models incrementally. Offsets between FEM and multibody nodes should be avoided unless strictly required.

Note about the FEM file: the format of the `fem_data_file` is relatively straightforward; it is reported in Appendix B.1. Initial modal displacements and velocities can be added, if required, by manually editing the file; however this practice is discouraged, unless strictly required.

Note about large FEM files: using a very large `fem_data_file` may require long time for reading and parsing ASCII floats. The keyword `use binary` instructs MBDyn to use a binary version of the `fem_data_file`. This binary version is automatically generated by MBDyn if requested by means of the keyword `create binary`. The binary version of the `fem_data_file` is used only if its timestamp is more recent than that of the ASCII version. The keyword `update binary` instructs MBDyn to regenerate the binary version when the ASCII version is more recent.

Note about the structural damping: structural damping can be provided by specifying any of the **no damping**, **proportional damping** or **diag damping**. In the last two cases, a damping factor is required. The corresponding damping for the i -th mode is computed as

$$c_i = \text{factor}_i 2\sqrt{k_i m_i}$$

where factor_i is the value provided for **proportional damping** (same for all modes) or the list of values provided for **diag damping** (for each mode). For example, a factor of 0.01 means 1% damping.

The keyword **diag damping** is followed by the count of damped modes, or by the keyword **all**. If **all** is used, as many damping coefficients as the available modes are expected. Otherwise, a list of **mode_damping** values follows. Each occurrence of **mode_damping** is made of the mode's index and the related damping value.

Output

Up to MBDyn 1.2.6 output occurred in a specific file that needed to be mandatorily given as the last argument to each modal element definition. Now output occurs in a `.mod` file, which contains, for each time step, as many rows as all the modes of all the modal elements whose output is enabled. Each row is structured as follows:

- a field containing the label of the modal joint and the number of the mode, separated by a dot (e.g. `label.mode`)
- the value of the modal unknown
- the value of the first derivative of the modal unknown
- the value of the second derivative of the modal unknown

Note: the number of the mode is the first field of the output is the ordinal of the mode in the FEM data file. If only a portion of the modes was selected using the **list** keyword, the mode numbers will reflect the selected modes. For example, if the **modal** joint was defined as

```
set: MODAL_ELEM = 99;
set: MODAL_NODE = 101;
joint: MODAL_ELEM, modal, MODAL_NODE,
      3, list, 1, 7, 9, # only use modes 1, 7 and 9 of a larger basis
      from file,
      "modal.fem",
      0;                # no fem nodes
```

the output will look like

```
99.1 0.000000e+00 0.000000e+00 0.000000e+00
99.7 0.000000e+00 0.000000e+00 0.000000e+00
99.9 0.000000e+00 0.000000e+00 0.000000e+00
```

Private Data

The following data is available, for each mode:

1. `"q[<m>]"` mode m value
2. `"qP[<m>]"` mode m derivative value

3. "qPP[<m>]" mode *m* second derivative value
4. "x[<n>,<i>]" FEM node *n* position *i* component value
5. "xP[<n>,<i>]" FEM node *n* velocity *i* component value
6. "xPP[<n>,<i>]" FEM node *n* acceleration *i* component value
7. "w[<n>,<i>]" FEM node *n* angular velocity *i* component value
8. "wP[<n>,<i>]" FEM node *n* angular acceleration *i* component value

When *m* is a number, the mode whose name is that number is used; when *m* is the character '#' followed by a number, the mode whose index is the number is used. The two cases may differ when only a subset of the modes defined in the FEM data file are actually used.

"FEM" node means that the label of the FEM node must be used; this allows to use the motion of nodes that are not attached to a multibody node. This motion is expressed in the global reference frame.

8.12.33 Plane displacement

This joint allows two nodes to move in the common relative 1–2 plane and to rotate about the common relative axis 3.

```
<joint_type> ::= plane displacement

<joint_arglist> ::=
  <node_1_label> ,
  position , (Vec3) <relative_offset_1>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label> ,
  position , (Vec3) <relative_offset_2>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]
```

Note: this element is temporarily disabled; combine an **in plane** and a **revolute rotation** (see Section 8.12.40) joint, or use the **total joint**, see Section 8.12.46.

8.12.34 Plane displacement pin

This joint allows a node to move in the relative 1–2 plane and to rotate about the relative axis 3 with respect to an absolute point and plane. See also Section 8.12.33.

```
<joint_type> ::= plane displacement pin

<joint_arglist> ::=
  <node_label> ,
  position , (Vec3) <relative_offset>
  [ , orientation , (OrientationMatrix) <relative_orientation_matrix> ] ,
  position , (Vec3) <absolute_pin_position>
  [ , orientation , (OrientationMatrix) <absolute_pin_orientation_matrix> ]
```

Note: this element is temporarily disabled; combine an **in plane** and a **revolute rotation** (see Section 8.12.40) joint instead, using a grounded node, or use a **total pin joint**, see Section 8.12.47.

8.12.35 Plane hinge

This joint has been renamed **revolute hinge** (see Section 8.12.38); the old name has been deprecated and its support may be discontinued in future versions.

8.12.36 Plane pin

This joint has been renamed **revolute pin** (see Section 8.12.39); the old name has been deprecated and its support may be discontinued in future versions.

8.12.37 Prismatic

This joints constraints the relative orientation of two nodes, so that their orientations remain parallel. The relative position is not constrained. The initial orientation of the joint must be compatible: use the **hinge** keyword to assign the joint initial orientation.

```
<joint_type> ::= prismatic

<joint_arglist> ::=
  <node_1_label>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ] ,
```

Hints

When wrapped by a **driven** element, the following hints are honored:

- **hinge{1}** the relative orientation of the joint with respect to node 1 is reset;
- **hinge{2}** the relative orientation of the joint with respect to node 2 is reset;

8.12.38 Revolute hinge

This joint only allows the relative rotation of two nodes about a given axis, which is axis 3 in the reference systems defined by the two **hinge** statements.

```
<joint_type> ::= revolute hinge

<joint_arglist> ::=
  <node_1_label> ,
    position , (Vec3) <relative_offset_1>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label> ,
    position , (Vec3) <relative_offset_2>
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]
  [ , initial theta , <initial_theta> ]
  [ , friction , <average_radius> ,
    [ preload , <const_value> , ]
    <friction_model> ,
    <shape_function> ]
```

Note: this element can be thought of as the combination of a **spherical hinge** (see Section 8.12.43), that constrains the three components of the relative position of the nodes, and of a **revolute rotation** (see Section 8.12.40), that constrains the relative orientation of the nodes so that only rotation about a common axis is allowed.

Rationale

The rationale for having two statements to indicate the position and orientation of the same entity is that the joint is supposed to constrain the position and orientation of two points, each attached to a node. This is what is typically known from the geometry of the components of a mechanism.

Output

The output occurs in the `.jnt` file, according to default joint output for the first 13 columns; specific columns contain:

- columns 14–16: the so-called *Euler* angles (in degrees) that describe the relative rotation; only the third component is relevant, the first two essentially indicate the accuracy of the rotation constraint;
- column 17–19: the relative angular velocity;

If friction is present:

- column 20: the friction moment about the revolute axis;
- subsequent columns: any friction model specific data.

Private Data

The following data is available:

1. **"rz"** relative rotation angle about revolute axis
2. **"wz"** relative angular velocity about revolute axis
3. **"Fx"** constraint reaction force in node 1 local direction 1
4. **"Fy"** constraint reaction force in node 1 local direction 2
5. **"Fz"** constraint reaction force in node 1 local direction 3
6. **"Mx"** constraint reaction moment about node 1 local direction 1
7. **"My"** constraint reaction moment about node 1 local direction 2
8. **"Mz"** constraint reaction moment about node 1 local direction 3

Hints

When wrapped by a **driven** element, the following hints are honored:

- **hinge{1}** the relative orientation of the joint with respect to node 1 is reset;
- **hinge{2}** the relative orientation of the joint with respect to node 2 is reset;
- **offset{1}** the offset of the joint with respect to node 1 is reset;
- **offset{2}** the offset of the joint with respect to node 2 is reset;
- unrecognized hints are passed through to the friction model, if any.

8.12.39 Revolute pin

This joint only allows the absolute rotation of a node about a given axis, which is axis 3 in the reference systems defined by the two `hinge` statements.

```
<joint_type> ::= revolute pin

<joint_arglist> ::=
  <node_label> ,
    position, (Vec3) <relative_offset>
    [ , orientation, (OrientationMatrix) <relative_orientation_matrix> ] ,
  position , (Vec3) <absolute_pin_position>
  [ , orientation , (OrientationMatrix) <absolute_pin_orientation_matrix> ]
  [ , initial_theta , <initial_theta> ]
```

Note: this is equivalent to a `revolute hinge` (see Section 8.12.38) when one node is grounded.

Private Data

The following data is available:

1. `"rz"` relative rotation angle about revolute axis
2. `"wz"` relative angular velocity about revolute axis
3. `"Fx"` constraint reaction force in node 1 local direction 1
4. `"Fy"` constraint reaction force in node 1 local direction 2
5. `"Fz"` constraint reaction force in node 1 local direction 3
6. `"Mx"` constraint reaction moment about node 1 local direction 1
7. `"My"` constraint reaction moment about node 1 local direction 2
8. `"Mz"` constraint reaction moment about node 1 local direction 3

Hints

When wrapped by a `driven` element, the following hints are honored:

- `hinge{1}` the relative orientation of the joint with respect to the node is reset;
- `hinge{0}` the relative orientation of the joint with respect to the absolute frame is reset;
- `offset{1}` the offset of the joint with respect to the node is reset;
- `offset{0}` the offset of the joint with respect to the absolute frame is reset;

8.12.40 Revolute rotation

This joint allows the relative rotation of two nodes about a given axis, which is axis 3 in the reference systems defined by the two `hinge` statements. The relative position is not constrained.

```
<joint_type> ::= revolute rotation

<joint_arglist> ::=
  <node_1_label>
    [ , position , (Vec3) <relative_offset_1> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_1> ] ,
  <node_2_label>
    [ , position , (Vec3) <relative_offset_2> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]
```

Rationale

A revolute joint without position constraints; this joint, in conjunction with an `inline` joint, should be used to constrain, for example, the two nodes of a hydraulic actuator.

Private Data

The following data is available:

1. `"rz"` relative rotation angle about revolute axis
2. `"wz"` relative angular velocity about revolute axis
3. `"Mx"` constraint reaction moment about node 1 local direction 1
4. `"My"` constraint reaction moment about node 1 local direction 2
5. `"Mz"` constraint reaction moment about node 1 local direction 3

Hints

When wrapped by a `driven` element, the following hints are honored:

- `hinge{1}` the relative orientation of the joint with respect to node 1 is reset;
- `hinge{2}` the relative orientation of the joint with respect to node 2 is reset;

8.12.41 Rod

The `rod` element represents a force between two nodes that depends on the relative position and velocity of two points, each rigidly attached to a `structural node`. The direction of the force is also based on the relative position of the points: it is the line that passes through them. If no offset is defined, the points are the nodes themselves.

The syntax is:

```
<joint_type> ::= rod

<joint_arglist> ::=
  <node_1_label> ,
```

```

    [ position , (Vec3) <relative_offset_1> , ]
<node_2_label> ,
    [ position , (Vec3) <relative_offset_2> , ]
{ (Scalar) <rod_length> | from nodes }
(ConstitutiveLaw<1D>) <const_law>

```

The constitutive law `const_law` receives as input the rod axial strain

$$\varepsilon = \frac{l - l_0}{l_0}, \quad (8.70)$$

where $l_0 = \text{rod_length}$ and l is the distance between the two points; in case of viscoelastic constitutive law, the axial strain derivative

$$\dot{\varepsilon} = \frac{\dot{l}}{l_0} \quad (8.71)$$

is also used. It returns the corresponding rod axial force,

$$F = F \left(\frac{l - l_0}{l_0}, \frac{\dot{l}}{l_0} \right).$$

If a `prestrain` ε_p is defined, it consists in an imposed value of the axial strain that is subtracted from the geometrical strain before being passed to the constitutive law; if a `prestress` F_p is defined, it consists in an imposed axial force value that is added to the one obtained from the constitutive law, namely

$$F = F_p + F \left(\frac{l - l_0}{l_0} - \varepsilon_p, \frac{\dot{l}}{l_0} \right).$$

For example, a linear elastic constitutive law for the rod element is defined as

$$F = EA \frac{l - l_0}{l_0},$$

where EA is the axial stiffness of the rod.

For further details on the supported constitutive laws, see Section 2.9.

Output

The output occurs in the `.jnt` file, which contains:

- the label
- the three components of the internal force in the reference frame of the element (column 2, component x , contains the force)
- the three components of the internal moment in the reference frame of the element (always zero)
- the three components of the internal force in the global reference frame
- the three components of the internal moment in the global reference frame (always zero)
- the length of the rod
- the three components of the rod axis in the global reference frame (a unit vector)
- the length rate of the rod (derivative of length wrt/ time)
- optional data appended by the constitutive law

Private Data

The following data is available:

1. **"F"** constraint force between the two nodes/connection points
2. **"L"** distance between the two nodes/connection points
3. **"LPrime"** distance rate between the two nodes/connection points

The rod joint can also access the private data provided by the constitutive law. It is accessed by prefixing the name of the data with the string **"constitutiveLaw."**; see the specific constitutive law description of the available data in Section 2.9.

8.12.42 Rod with offset

The syntax is:

```
<joint_type> ::= rod with offset

<joint_arglist> ::=
  <node_1_label> ,
    (Vec3) <relative_offset_1> ,
  <node_2_label> ,
    (Vec3) <relative_offset_2> ,
  { (Scalar) <rod_length> | from nodes } ,
  (ConstitutiveLaw<1D>) <const_law>
```

Analogous to the **rod** joint with the optional offsets; see Section 8.12.41 for details.

8.12.43 Spherical hinge

This joint constrains the relative position of two nodes; the relative orientation is not constrained.

```
<joint_type> ::= spherical hinge

<joint_arglist> ::=
  <node_1_label> ,
    [ position , (Vec3) <relative_offset_1> , ]
    [ orientation , (OrientationMatrix) <relative_orientation_matrix_1> , ]
  <node_2_label>
    [ , position , (Vec3) <relative_offset_2> ]
    [ , orientation , (OrientationMatrix) <relative_orientation_matrix_2> ]
```

Note: the orientation matrix, set by means of the **orientation** keyword, is used for output purposes only. A default identity matrix is assumed.

Hints

When wrapped by a **driven** element, the following hints are honored:

- **offset{1}** the offset of the joint with respect to node 1 is reset;
- **offset{2}** the offset of the joint with respect to node 2 is reset;

8.12.44 Spherical pin

This joint constrains the absolute position of a node; the relative orientation is not constrained.

```
<joint_type> ::= spherical pin

<joint_arglist> ::=
  <node_label> ,
  [ position , (Vec3) <relative_offset> , ]
  [ orientation , (OrientationMatrix) <relative_orientation> , ]
  position , (Vec3) <absolute_pin_position>
  [ , orientation , <absolute_orientation> ]
```

Note: this is equivalent to a **spherical hinge** (see Section 8.12.43) when one node is grounded. An alternative way to model a grounded spherical hinge requires the use of another node, clamped by a **clamp** joint.

8.12.45 Total equation

Original implementation: Alessandro Fumagalli, Marco Morandini.

See also **total reaction** (see Section 8.12.48). TODO

8.12.46 Total joint

Original implementation: Alessandro Fumagalli;

Review: Pierangelo Masarati.

This element allows to arbitrarily constrain specific components of the relative position and orientation of two nodes. The value of the constrained components of the relative position and orientation can be imposed by means of drives. As such, this element allows to mimic the behavior of most ideal constraints that connect two nodes.

```
<joint_type> ::= total joint

<joint_arglist> ::=
  <node_1_label>
  [ , position , (Vec3) <relative_offset_1> ]
  [ , position orientation , (OrientationMatrix) <rel_pos_orientation_1> ]
  [ , rotation orientation , (OrientationMatrix) <rel_rot_orientation_1> ]
  , <node_2_label>
  [ , position , (Vec3) <relative_offset_2> ]
  [ , position orientation , (OrientationMatrix) <rel_pos_orientation_2> ]
  [ , rotation orientation , (OrientationMatrix) <rel_rot_orientation_2> ]
  [ , position constraint ,
    <position_status> , <position_status> , <position_status> ,
    (TplDriveCaller<Vec3>) <imposed_relative_position> ]
  [ , orientation constraint ,
    <orientation_status> , <orientation_status> , <orientation_status> ,
    (TplDriveCaller<Vec3>) <imposed_relative_rotation> ]

<position_status> ::=
  { inactive | active | position | velocity | <status> }
```

```

<orientation_status> ::=
    { inactive | active | rotation | angular velocity | <status> }

```

The relative position imposed by the **position constraint** is imposed in a reference frame rigidly attached to the first node, in the optional offset **relative_offset_1**, and optionally further oriented by the **rel_pos_orientation_1** matrix.

The relative orientation imposed by the **orientation constraint** is imposed in a reference frame rigidly attached to the first node, optionally further oriented by the **rel_rot_orientation_1** matrix. It consists in the Euler vector that expresses the imposed relative orientation, in radian.

The keyword **active** means that the constraint is active with respect to that component of relative motion, so the related motion component is constrained, while **inactive** means that the constraint is not active with respect of that component of relative motion, so the related motion component is not constrained. Otherwise, a boolean can be provided in **status** to indicate that the degree of constraint is either inactive (0) or active (1); this may be useful, for instance, to make constraint activation conditional in parametric input files. The same applies to the status of the components of the impose orientation.

If a component of relative position or orientation is active, the corresponding component of the imposed position or orientation is enforced, otherwise it is ignored; however, the complete three-dimensional vectors of the imposed relative position or orientation must be provided.

When the **position constraint** is enforced, the keyword **active** is equivalent to **position**; if the keyword **velocity** is used, the constraint imposes the velocity of that component of the relative position, resulting in a non-holonomic constraint.

Similarly, when the **orientation constraint** is enforced, the keyword **active** is equivalent to **rotation**; if the keyword **angular velocity** is used, the constraint imposed the angular velocity of that component of the relative orientation, resulting in a non-holonomic constraint.

Output

The output occurs in the **.jnt** file, which contains:

- column 1: the label
- columns 2–4: the three components of the internal force in the reference frame of the element
- columns 5–7: the three components of the internal moment in the reference frame of the element
- columns 8–10: the three components of the internal force in the global reference frame
- columns 11–13: the three components of the internal moment in the global reference frame
- columns 14–16: the three components of the relative displacement in the reference frame of the element
- columns 17–19: the three components of the relative rotation vector in the reference frame of the element
- columns 20–22: the three components of the relative velocity in the reference frame of the element
- columns 23–25: the three components of the relative angular velocity in the reference frame of the element

Private Data

The following data is available:

1. "px" relative position along node 1 position orientation axis 1
2. "py" relative position along node 1 position orientation axis 2
3. "pz" relative position along node 1 position orientation axis 3
4. "rx" relative orientation about node 1 rotation orientation axis 1
5. "ry" relative orientation about node 1 rotation orientation axis 2
6. "rz" relative orientation about node 1 rotation orientation axis 3
7. "Fx" reaction force about node 1 position orientation axis 1
8. "Fy" reaction force about node 1 position orientation axis 2
9. "Fz" reaction force about node 1 position orientation axis 3
10. "Mx" reaction moment about node 1 rotation orientation axis 1
11. "My" reaction moment about node 1 rotation orientation axis 2
12. "Mz" reaction moment about node 1 rotation orientation axis 3
13. "dx" imposed relative position along node 1 position orientation axis 1
14. "dy" imposed relative position along node 1 position orientation axis 2
15. "dz" imposed relative position along node 1 position orientation axis 3
16. "tx" imposed relative orientation about node 1 rotation orientation axis 1
17. "ty" imposed relative orientation about node 1 rotation orientation axis 2
18. "tz" imposed relative orientation about node 1 rotation orientation axis 3
19. "vx" relative velocity along node 1 rotation orientation axis 1
20. "vy" relative velocity along node 1 rotation orientation axis 2
21. "vz" relative velocity along node 1 rotation orientation axis 3
22. "wx" relative angular velocity about node 1 rotation orientation axis 1
23. "wy" relative angular velocity about node 1 rotation orientation axis 2
24. "wz" relative angular velocity about node 1 rotation orientation axis 3

Hints

When wrapped by a `driven` element, the following hints are honored:

- `offset{1}` the offset of the joint with respect to node 1 is reset;
- `offset{2}` the offset of the joint with respect to node 2 is reset;
- `position-hinge{1}` the relative orientation of the relative position constraint with respect to node 1 is reset;
- `position-hinge{2}` the relative orientation of the relative position constraint with respect to node 2 is reset;
- `orientation-hinge{1}` the relative orientation of the relative orientation constraint with respect to node 1 is reset;
- `orientation-hinge{2}` the relative orientation of the relative orientation constraint with respect to node 2 is reset;
- `position-drive3` resets the relative position drive; the hint is passed to the Vec3 drive hint parser.
- `orientation-drive3` resets the relative orientation drive; the hint is passed to the Vec3 drive hint parser.

Example. This example mimicks an instance of NASTRAN's RBE2 element that constrains displacement in x and y , and rotation about z :

```
set: const integer EID = 100;
set: const integer GN = 200;
set: const integer GM1 = 300;

# ...

# $......2.....3.....4.....5.....
# RBE2   EID   GN   CM   GM1
joint: EID, total joint,
      GN,
      position, reference, node, null,
      position orientation, reference, node, eye,
      rotation orientation, reference, node, eye,
GM1,
      position, reference, other node, null,
      position orientation, reference, other node, eye,
      rotation orientation, reference, other node, eye,
position constraint,
      active, # "active" if CM includes "1"; "inactive" otherwise
      active, # "active" if CM includes "2"; "inactive" otherwise
      inactive, # "active" if CM includes "3"; "inactive" otherwise
      null,
orientation constraint,
      inactive, # "active" if CM includes "4"; "inactive" otherwise
      inactive, # "active" if CM includes "5"; "inactive" otherwise
      active, # "active" if CM includes "6"; "inactive" otherwise
      null;
```

8.12.47 Total pin joint

This element allows to arbitrarily constrain specific components of the absolute position and orientation of a node. The value of the constrained components of the absolute position and orientation can be imposed by means of drives. As such, this element allows to mimic the behavior of most ideal constraints that ground one node.

```
<joint_type> ::= total pin joint

<joint_arglist> ::=
  <node_label>
  [ , position , (Vec3) <relative_offset> ]
  [ , position orientation , (OrientationMatrix) <rel_pos_orientation> ]
  [ , rotation orientation , (OrientationMatrix) <rel_rot_orientation> ]
  [ , position , (Vec3) <absolute position> ]
  [ , position orientation , (OrientationMatrix) <abs_pos_orientation> ]
  [ , rotation orientation , (OrientationMatrix) <abs_rot_orientation> ]
  [ , position constraint ,
    <position_status> , <position_status> , <position_status> ,
    (TplDriveCaller<Vec3>) <imposed_absolute_position> ]
  [ , orientation constraint ,
    <orientation_status> , <orientation_status> , <orientation_status> ,
    (TplDriveCaller<Vec3>) <imposed_absolute_rotation> ]

<position_status> ::=
  { inactive | active | position | velocity | <status> }

<orientation_status> ::=
  { inactive | active | rotation | angular velocity | <status> }
```

The non-holonomic variant is not implemented yet.

Output

The output occurs in the .jnt file, which contains:

- column 1: the label
- columns 2–4: the three components of the internal force in the reference frame of the element
- columns 5–7: the three components of the internal moment in the reference frame of the element
- columns 8–10: the three components of the internal force in the global reference frame
- columns 11–13: the three components of the internal moment in the global reference frame
- columns 14–16: the three components of the absolute displacement in the reference frame of the element
- columns 17–19: the three components of the absolute rotation vector in the reference frame of the element
- columns 20–22: the three components of the absolute velocity in the reference frame of the element
- columns 23–25: the three components of the absolute angular velocity in the reference frame of the element

Private Data

The following data is available:

1. "px" absolute position along absolute position orientation axis 1
2. "py" absolute position along absolute position orientation axis 2
3. "pz" absolute position along absolute position orientation axis 3
4. "rx" absolute orientation about absolute rotation orientation axis 1
5. "ry" absolute orientation about absolute rotation orientation axis 2
6. "rz" absolute orientation about absolute rotation orientation axis 3
7. "Fx" reaction force about absolute position orientation axis 1
8. "Fy" reaction force about absolute position orientation axis 2
9. "Fz" reaction force about absolute position orientation axis 3
10. "Mx" reaction moment about absolute rotation orientation axis 1
11. "My" reaction moment about absolute rotation orientation axis 2
12. "Mz" reaction moment about absolute rotation orientation axis 3
13. "dx" imposed absolute position along absolute position orientation axis 1
14. "dy" imposed absolute position along absolute position orientation axis 2
15. "dz" imposed absolute position along absolute position orientation axis 3
16. "tx" imposed absolute orientation about absolute rotation orientation axis 1
17. "ty" imposed absolute orientation about absolute rotation orientation axis 2
18. "tz" imposed absolute orientation about absolute rotation orientation axis 3
19. "vx" absolute velocity along absolute position orientation axis 1
20. "vy" absolute velocity along absolute position orientation axis 2
21. "vz" absolute velocity along absolute position orientation axis 3
22. "wx" absolute angular velocity about absolute rotation orientation axis 1
23. "wy" absolute angular velocity about absolute rotation orientation axis 2
24. "wz" absolute angular velocity about absolute rotation orientation axis 3

Hints

When wrapped by a `driven` element, the following hints are honored:

- `offset{1}` the offset of the joint with respect to the node is reset;
- `offset{0}` the absolute position of the joint is reset;
- `position-hinge{1}` the relative orientation of the absolute position constraint with respect to the node is reset;
- `position-hinge{0}` the absolute orientation of the absolute position constraint is reset;
- `orientation-hinge{1}` the relative orientation of the absolute orientation constraint with respect to the node is reset;
- `orientation-hinge{0}` the absolute orientation of the absolute orientation constraint is reset;
- `position-drive3` resets the absolute position drive; the hint is passed to the Vec3 drive hint parser.
- `orientation-drive3` resets the absolute orientation drive; the hint is passed to the Vec3 drive hint parser.

8.12.48 Total reaction

Original implementation: Alessandro Fumagalli, Marco Morandini.

See also `total equation` (see Section 8.12.45). TODO

8.12.49 Universal hinge

Deprecated in favour of the `cardano hinge` (see Section 8.12.6).

8.12.50 Universal pin

Deprecated in favour of the `cardano pin` (see Section 8.12.7).

8.12.51 Universal rotation

Deprecated in favour of the `cardano rotation` (see Section 8.12.8).

8.12.52 Viscous body

```
<joint_type> ::= viscous body

<joint_arglist> ::=
  <node_label> ,
  [ position , (Vec3) <relative_offset> , ]
  [ orientation , (OrientationMatrix) <relative_orientation> , ]
  (ConstitutiveLaw<6D>) <const_law>
```

This element defines a force and a moment that depend on the absolute linear and angular velocity of a body, projected in the reference frame of the node itself. The force and moment are defined as a 6D viscous constitutive law.

8.12.53 Lower Pairs

Revolute

Use a **revolute hinge**, or a **total joint**.

Prismatic

Use a **prismatic** and an **inline**, or a **total joint**.

Screw

Not supported yet.

Cylindrical

Use a **revolute rotation** and an **inline**, or a **total joint**.

Spherical

Use a **spherical hinge**, or a **total joint**.

Planar

Use a **revolute rotation** and an **in plane**, or a **total joint**.

8.13 Joint Regularization

The **joint regularization** element is used to modify algebraic constraint equations in order to improve ill-conditioned problems.

```
<elem_type> ::= joint regularization  
  
<normal_arglist> ::= <type> [ , <data> ]  
  
<type> ::= tikhonov  
  
<data> ::= { <coef> | list , <coef_1> [ , ... ] }
```

An element is instantiated, which requires an underlying algebraic constraint (a **joint** element that instantiates algebraic equations) with the same label to exist. In that case, the algebraic constraint equations are modified to regularize the problem. The **joint** that is regularized must exist, and it must write constraint equations. If the keyword **list** is used, the coefficients for each internal state are expected; the number of coefficients is determined by the joint.

8.13.1 Tikhonov

The **tikhonov** joint regularization type consists in modifying a constraint, expressed by an algebraic equation $\Phi(\mathbf{x}, t)$ and the corresponding Lagrangian multipliers $\boldsymbol{\lambda}$, namely

$$\left(\Phi_{/x}^T \boldsymbol{\lambda}\right)_{/x} \Delta \mathbf{x} + \Phi_{/x}^T \Delta \boldsymbol{\lambda} = \mathbf{F} - \Phi_{/x}^T \boldsymbol{\lambda} \quad (8.72)$$

$$\Phi_{/x} \Delta \mathbf{x} = -\Phi(\mathbf{x}, t) \quad (8.73)$$

by adding the multipliers to the constraint equation, weighed by the coefficient `coef` (or by the list of coefficients `coef_<i>`, in case different coefficients are given):

$$\left(\Phi_{/x}^T \lambda \right)_{/x} \Delta x + \Phi_{/x}^T \Delta \lambda = F - \Phi_{/x}^T \lambda \quad (8.74)$$

$$\Phi_{/x} \Delta x - \langle \text{coef} \rangle \cdot \Delta \lambda = -\Phi(x, t) + \langle \text{coef} \rangle \cdot \lambda \quad (8.75)$$

The Tikhonov regularization allows the constraint to be violated by an amount that depends on the multipliers. In this sense, the coefficient `coef` should be considered sort of a compliance: the larger the coefficient, the larger the constraint violation for a given value of the reaction λ .

8.14 Plate Elements

8.14.1 Shell4

Authors: Marco Morandini and Riccardo Vescovini

The shell4 elements model a four-node shell. The syntax is

```
<elem_type> ::= { shell4eas | shell4easans }

<normal_arglist> ::=
  <node_1_label> , <node_2_label> , <node_3_label> , <node_4_label> ,
  <shell_constitutive_law_data>

<shell_constitutive_law_data> ::=
  { [ matr , ] <12x12_matrix>
    | sym , <upper_triangular_12x12_matrix>
    | diag , <diagonal_12x12_matrix>
    | isotropic , <isotropic_data> }
  [ , prestress , (Vec12) <prestress> ]

<isotropic_data> ::=
  { { E | Young modulus } , <E>
    | { nu | Poisson modulus } , <nu>
    | { G | shear modulus } , <G>
    | as , <as> # TODO: clarify
    | at , <at> # TODO: clarify
    | thickness , <thickness> }
  [ , ... ]
```

The names `shell4eas` and `shell4easans` respectively stand for “Enhanced Assumed Strain” (EAS) and “Enhanced Assumed Strain-Assumed Natural Strain” (EAS-ANS).

Nodes are numbered according to Figure 8.8.

Only linear elastic constitutive properties can be currently modeled. They consist in a 12×12 matrix




Figure 8.8: Shell: definitions

that expresses the force and moment fluxes as functions of linear and angular strains according to

$$\begin{pmatrix} n_{11} \\ n_{12} \\ n_{13} \\ n_{21} \\ n_{22} \\ n_{23} \\ m_{11} \\ m_{12} \\ m_{13} \\ m_{21} \\ m_{22} \\ m_{23} \end{pmatrix} = [D] \begin{pmatrix} \varepsilon_{11} \\ \varepsilon_{12} \\ \varepsilon_{13} \\ \varepsilon_{21} \\ \varepsilon_{22} \\ \varepsilon_{23} \\ \kappa_{11} \\ \kappa_{12} \\ \kappa_{13} \\ \kappa_{21} \\ \kappa_{22} \\ \kappa_{23} \end{pmatrix} \quad (8.76)$$

When the `isotropic` keyword is used, only two of the optional sub-keywords `E`, `nu` and `G` are required, as the remaining parameter can be computed from the other two according to the relationship

$$\mathbf{G} = \frac{\mathbf{E}}{2(1 + \nu)}. \quad (8.77)$$

If all are provided, they must be consistent. The optional parameters `as` and `at` are not documented yet; the default should be used.

TODO: improve

8.15 User-Defined Elements

8.15.1 Loadable Element

Note: the `loadable` element is deprecated in favor of the `user defined` element.

The `loadable` element is a wrapper for a user-defined element that is compiled in a separated module and linked run-time. The module should provide a comprehensive set of functions according to a specified API; default functions are available if no special features are required. Implementation of modules can be very easy, but a deep knowledge of the internals of the code is required when special tasks need to be performed. There are virtually no limits on what a loadable element can do.

The syntax is simply:

```

<elem_type> ::= loadable

<normal_arglist> ::= " <module_name> "
    [ , name , " <calls> " ]
    [ , <module_data> ]

```

where `module_name` is the name of the module file; as soon as the file is checked and the binding of the structure with function bindings succeeded, a function called `read()` is invoked, passing it the input stream. This function is in charge of reading `module_data` following the general syntax of the input file.

An alternative form is

```

<normal_arglist> ::= reference , " <name> "
    [ , <module_data> ]

```

where `name` is the name by which the loadable element recorded itself when registered via the `module load` directive, as described in Section 2.3.6. As a consequence, the following forms are equivalent:

```

# direct runtime loading
loadable: 1, "/some/module.so";
# "/some/module.so" registers itself as "some_module"
module load: "/some/module.so";
loadable: 2, reference, "some_module";
# works also as joint (might obsolete loadable elements)
joint: 3, some_module;

```

It is advisable that the function `read()` prints some help message when the first field of `module_data` is the keyword `help`. All the helpers and the high-level structures are available, such as drivers, constitutive laws, reference frames. Refer to each module for a description (if available) of the features and of the input/output format. `module_name` should be a full path to the module function. If the name starts with a slash “/”, the full path name is used. Otherwise the module is searched in the colon-separated list of directories contained in the environment variable `LD_LIBRARY`, then among the libraries listed in `/etc/ld.so.cache`, and finally in `/usr/lib` and in `/lib` (see `dlopen(3)`). At last, it is searched in the current directory, and the extension `.so` is added if missing. The string `calls` represents the name of the structure that contains the bindings to the functions. The default is `calls`.

Refer to `$(BASE)/mbdyn/base/loadable.h` for a description of the functions that are allowed. An example module is given in directory

```

$(BASE)/modules/module-template/

```

which can be used as a starting point to build a custom module. The `loadable` element interface allows to link modules in different languages, e.g. C or FORTRAN77; simply use `module-template` as a guideline to providing callbacks to the `loadable` element interface and to collect required info from the main program (e.g. node positions, equation indices and everything else that is required for appropriate connection), then call the functions that actually do the work in other languages from inside the callbacks.

8.15.2 User-Defined Element

The `user defined` element is a much more streamlined form of custom element definition than the `loadable` element. From the point of view of the syntax the differences are minimal; however, from an implementation point of view, the `user defined` element is preferable.

The definition of a `user defined` element requires two steps. In the first step, a run-time loadable module is loaded using the `module load` directive (see Section 2.3.6). This registers a handler to the

user defined element type along with a **name** used to reference it. In the second step, an instance of that **user defined** element type is created, referencing it by **name**.

The syntax of the **user defined** element is

```
<elem_type> ::= user defined  
  
<normal_arglist> ::= <name> [ , <module_data> ]
```

As for the **loadable** element, it is recommended that some useful output is given if the first **module_data** is the keyword **help**.

An example module is given in directory

```
$(BASE)/modules/module-template2/
```

which can be used as a starting point to build a custom module.

Example.

```
module load: "libmodule-template2"  
  
user defined: 1000, template2, help;
```

8.15.3 General Discussion on Run-Time Loadable Modules

In general, to call external functions from C++ one needs to declare them as

```
#include <sys/types.h>  
extern "C" {  
    int a_C_function(int arg, double darg);  
    int a_F77_subroutine(int32_t *iarg, double *darg);  
}
```

The same applies to FORTRAN 77 functions; only, the naming convention usually is compiler dependent; some compilers turn all the names to uppercase or lowercase (remember that FORTRAN 77 is case insensitive); other compilers add underscores at the beginning or at the end of the names. Check what is the behavior of your compiler, by compiling a simple program with your favorite FORTRAN 77 compiler, then looking at it with the utility `nm(1)`, which will show how the symbols are represented internally. For instance, the code

```
C This is a compiler test  
  SUBROUTINE F77SUB(I, D)  
    INTEGER*4 I  
    REAL*8 D(*)  
  
    D(I) = 0.0  
  
  END
```

when compiled with `g77(1)` on a GNU/Linux system, yields:

```
[masarati@mbdyn manual]$ g77 -c f77sub.f  
[masarati@mbdyn manual]$ nm f77sub.o  
00000000 T f77sub_
```

That is, `g77(1)` lowercases all symbols, and adds a trailing underscore. Macros to automatically detect and take care of this behavior are planned.

To compile loadable modules, one needs to configure the package as follows:

```
./configure --with-module=<module_name>
```

where `module_name` is the name of the directory the module is placed in with the `module-` part stripped; e.g. to compile the `tire` module that resides in `$(BASE)/modules/module-wheel2` one must type

```
./configure --with-module=wheel2
```

Multiple modules can be compiled by typing the list of the names separated by blanks.

The modules need to resolve some of the symbols that are in the main executable; until a full working libtool support is implemented, this must be done by hand. The `g++(1)` compiler requires the switch `'-rdynamic'` to be added to the loader's flags.

For example,

```
./configure --with-module=<module_name> LDFLAGS="-rdynamic"
```

8.16 Output Elements

Output elements take care of inter-process communication. These elements can use specific communication means, depending on the type of simulation they are used for, and can communicate specific types of data.

8.16.1 Stream output

This is a special element which takes care of sending output to external processes by means of either `local` or `inet` sockets during batch or real-time simulations. This topic is under development, so expect frequent changes, and please do not count too much on backward compatibility.

The syntax is:

```
<elem_type> ::= stream output

<arglist> ::=
  stream name , " <stream_name> " ,
  create , { yes | no } ,
  [ { local , " <socket_name> " , |
    [ port , <port_number> , ]
    [ host , " <host_name> " , ] } ]
  [ { [ no ] signal
    | [ non ] blocking
    | [ no ] send first
    | [ do not ] abort if broken } [ , ... ] , ]
  [ output every , <steps> , ]
  [ echo , <file_name>
    [ , precision , <precision> ]
    [ , shift , <shift> ] , ]
  <content>
```

The stream output allows MBDyn to send streamed outputs to remote processes during both batch and real-time simulations, using sockets either in the `local` or in the `inet` namespace. If the simulation is run in real-time using RTAI, RTAI mailboxes are used instead.

- `stream_name` is the name of the RTAI mailbox where the output is written (a unique string no more than six characters long); it is basically ignored by the `stream output` element except when using RTAI;
- the `create` keyword determines whether the socket must be created or looked-up as already existing on the system; if `create` is set to `no`, MBDyn will retry for 60 seconds and then give up;
- `socket_name` is the path of the `local` socket that is used to communicate between processes;
- `port_number` is the port number to be used with a `inet` socket. The default port number is 9011 (intentionally unassigned by IANA). If no `host_name` is given, `localhost` will be used;
- `host_name` is the name or the IP of the remote host where the mailbox resides; note that if this field is given, `create` must be set to `no`. The simulation will not start until the socket is created on the remote host;
- the flag `no signal` requests that no SIGPIPE be raised when sending through a socket when the other end is broken (by default, SIGPIPE is raised);
- the flag `non blocking` requests that operations on the socket do not block (or block, in case of `blocking`, the default);
- the flag `no send first` requests that no send occurs before the first time step (by default, data is always sent);
- the flag `do not abort if broken` requests that the simulation continues in case the connection breaks. No further data send will occur for the duration of the simulation (the default);
- the field `output every` requests output to occur only every `steps`;
- the field `echo` causes the content that is sent to the peer to be echoed on file `file_name`; the optional parameter `precision` determines the number of digits used in the output; the optional parameter `shift` is currently unused;
- the field `content` is detailed in the next section.

This element, when used with the `motion` content type, obsoletes the `stream motion output` element (see Section 8.16.3). When the simulation is executed in real-time using RTAI, this element obsoletes the `RTAI output` element (see Section 8.16.2).

Streamed output

Different types of data can be sent. The most general form is called `values`, consisting in an arbitrary set of independent scalar channels. A form specifically intended to communicate the motion of a mechanical system is called `motion`, consisting in a subset of the kinematics of a set of structural nodes:

```
<content> ::= { <values> | <motion> }
```

```
<values> ::= [ values , ]
           <channel_number> ,
           <value>
```

```

    [ , ... ]

<value> ::=
    { [ nodedof , ] (NodeDof) <output_dof>
      | drive , (DriveCaller) <drive_data> }

<motion> ::= motion ,
    [ output flags ,
      { position
        | orientation matrix
        | orientation matrix transpose
        | velocity
        | angular velocity }
      [ , ... ] , ]
    { all | <struct_node_label> [ , ... ] }

```

where

- the (optional) keyword **values** indicates that a set of independent scalar channels is output by the element;
- the number of channels **channel_number** that are written determines how many **value** entries must be read. In case of **nodedof** (the default, deprecated), they must be valid scalar dof entries, which can be connected in many ways to nodal degrees of freedom; in case of **drive**, they can be arbitrary functions, including node or element private data;
- the keyword **motion** indicates that a subset of the kinematic parameters of a set of structural nodes is output by the element. As opposed to the **values** case, which is intended for generic interprocess communication output, this content type is intended to ease and optimize the output of the motion of structural nodes, to be used for on-line visualization purposes. By default, only the position of the selected nodes is sent. This is intended for interoperability with a development version of EasyAnim which can read the motion info (the so-called “van” file) from a stream. The optional keyword **output flags** allows to request the output of specific node kinematics: the node position, orientation matrix (either row- or column-major), the velocity and the angular velocity. The default is equivalent to **output flags, position**.

Non real-time simulation

During non real-time simulations, streams operate in blocking mode. The meaning of the parameters is:

- **stream_name** indicates the name the stream would be known as by RTAI; it must be no more than 6 characters long, and mostly useless;
- the instruction **create** determines whether MBDyn will create the socket, or try to connect to an existing one;
- the keyword **local** indicates that a socket in the local namespace will be used; if **create** is set to **yes**, the socket is created, otherwise it must exist.
- either of the keywords **port** or **host** indicate that a socket in the internet namespace will be used; if **create** is set to **yes**, **host_name** indicates the host that is allowed to connect to the socket; it defaults to any host (0.0.0.0); if **create** is set to **no**, **host_name** indicates what host to connect to; the default is localhost (127.0.0.1); the default port is **9011** (intentionally unassigned by IANA);

- the flag `no signal` is honored;
- the flag `non blocking` is honored;
- the flag `no send first` is honored;
- the flag `do not abort if broken` is honored.

If no socket type is specified, i.e. none of the `local`, `port` and `host` keywords are given, a socket is opened by default in the internet namespace with the default IP and port; the `create` keyword is mandatory.

Real-time simulation

During real-time simulations, streams wrap non-blocking RTAI mailboxes. The meaning of the parameters is:

- the parameter `stream_name` indicates the name the stream will be known as in RTAI's resource namespace; it must be exactly 6 characters long;
- the instruction `create` determines whether the mailbox will be created or looked for by MBDyn;
- the keyword `local` is ignored;
- the keyword `host` indicates that a mailbox on a remote host will be used; it is useless when `create` is set to `yes`, because RTAI does not provide the possibility to create remote resources; if none is given, a local mailbox is assumed;
- the keyword `port` is ignored.
- the flag `no signal` is ignored;
- the flag `non blocking` is honored; however, blocking mailboxes make little sense, and real-time synchronization using RTAI should not rely on blocking mailboxes;
- the flag `no send first` is ignored (although it should be honored when the mailbox is blocking);
- the flag `do not abort if broken` is ignored; the program is always terminated if a mailbox is no longer available.

8.16.2 RTAI output

This element is actually used only when the simulation is scheduled using RTAI; otherwise, the corresponding `stream output` element is used (see Section 8.16.1). As a consequence, its explicit use is discouraged and deprecated. The `stream output` element should be used instead.

8.16.3 Stream motion output

This element type is obsoleted by the `stream output` element with the `motion` content type (see Section 8.16.3). The syntax is:

```
<elem_type> ::= stream motion output

<arglist> ::=
    stream name , " <stream_name> " ,
    create , { yes | no } ,
```

```

[ { local , " <socket_name> " ,
  | [ port , <port_number> , ] [ host , " <host_name> " , ] } ]
[ { [ no ] signal
  | [ non ] blocking
  | [ no ] send first
  | [ do not ] abort if broken }
  [ , ... ] , ]
<motion>

```

Its support may be discontinued in the future.

8.17 Miscellaneous

This section lists some extra cards that do not correspond to any specific simulation entity, but rather alter the behavior of existing entries or cause special operations to be undertaken during model input.

8.17.1 Bind

The statement `bind` does not really define an element. It is rather used to instruct a `parameter node` about which parameter of an element it is bound to. The `parameter node` must exist, and the element the node is being bound to, of type `elem_type` and label `element_label`, must have been already defined. The complete syntax is:

```

<elem_type> ::= bind

<arglist> ::=
  <element_label> ,
  <element_type> ,
  <parameter_node_label> ,
  <bind_args>

```

where `bind_args` depend on the type of parameter node.

Element

When binding an element to an `element` parameter node, each element makes a number of specific parameters available. A detailed list is given for each element in the private data section. In that case:

```

<bind_args> ::=
  { [ index , ] <parameter_index>
    | string , " <parameter_name> " }

```

The value of `parameter_index` must be legal, i.e. between 1 and the maximum number of parameters made available by the element. The alternative form, using the keyword `string` followed by the `parameter_name`, allows more friendly definition of the binding. The name of the parameter depends on the element whose property is being bound. A complete listing of the parameters that a parameter node can be bound to can be found in the ‘Private data’ subsection of each element’s specification.

Example. The parameter node `ANGLE` is bound to the rotation of a `revolute hinge` (see Section 8.12.38).

```

# ... problem block

begin: control data;
  structural nodes: 2;
  parameter nodes: 1;
  forces: 2;
  # ... other control data
end: control data;

set: integer NODE1 = 1000;
set: integer NODE2 = 2000;
set: integer ANGLE = 5000;
set: integer REVOLUTE = 6000;

begin: nodes;
  structural: NODE1, dynamic, null, eye, null, null;
  structural: NODE2, dynamic, null, eye, null, null;
  parameter: ANGLE, element;

  # ... other nodes
end: nodes;

begin: elements;
  joint: REVOLUTE, revolute hinge,
    NODE1,
    position, reference, node, null,
    orientation, reference, node, eye,
    NODE2,
    position, reference, node, null,
    orientation, reference, node, eye;
  bind: REVOLUTE, joint, ANGLE, string, "rx";
  couple: 1, NODE1, 0.,0.,1.,
    dof, ANGLE, parameter, 1, linear, 0.,1.;
  couple: 2, NODE2, 0.,0.,1.,
    element, REVOLUTE, joint, string, "rx", linear, 0.,1.;

  # ... other elements
end: elements;

```

Note that the same element data, i.e. the revolute hinge relative rotation angle, is used to drive a couple in two different ways; the latter, by means of the **element** drive (see Section 2.5.11) is more direct, but the former, by means of the **dof** drive (see Section 2.5.7) through the **bind** mechanism has the additional effect of updating the **parameter** node, which can be used to connect **genel** elements for special purposes.

Beam strain gage

When binding an element to a **beam strain gage** parameter node, the **element_type** field must be **beam**. In that case:

```
<bind_args> ::= <beam_evaluation_point>
```

where `beam_evaluation_point` is the evaluation point of the beam element where the internal strain and curvatures must be evaluated. It must be 1 for 2-node beams, while it can be either 1 or 2 for 3-node beams.

Example.

```
# ... problem block

begin: control data;
  parameter nodes: 1;
  beams: 1;
  # ... other control data
end: control data;

set: integer BEAM = 100;
set: integer STRAIN = 200;

begin: nodes;
  parameter: STRAIN, beam strain gage, 0.0, 0.1;
  # ... other nodes
end: nodes;

begin: elements;
  beam3: BEAM, ...; # beam data
  bind: BEAM, beam, STRAIN, 1;
  # ... other elements
end: elements;
```

8.17.2 Driven Element

The `driven` type is not an element by itself. It is a wrapper that masks another element and switches it on and off depending on the (boolean) value of a drive. It can be used to emulate a variable topology model, where some elements simply don't contribute to the residual or to the Jacobian matrices when their drive has a certain value. Since the drivers can be arbitrary functions of the time, or other parameters including the value of any degree of freedom, the driven elements can be "driven" in a very flexible way. Every element can be driven, except those that can be instantiated once only. The syntax for a driven element is:

```
<elem_type> ::= driven

<normal_arglist> ::= (DriveCaller) <element_driver> ,
  [ hint , " <hint> " [ , ... ] ] ,
  <driven_element>

<driven_element> ::=
  { existing : <driven_elem_type> , <driven_elem_label>
    | <element_card> }
```

When the keyword `existing` is used, an existing element of type `driven_elem_type` and label `driven_elem_label` is looked for, and it is wrapped by the driven element. In this case, no new element is instantiated. The label of the element must match that of the driving element given at the beginning. For consistency with

the syntax, and for more flexibility, even when wrapping an existing element the output flags can be set at the end of the card. This flag overrides the one set when the driven element was instantiated.

Otherwise, a regular element is read after the driving element's declaration; it is then instantiated and wrapped by the **driven** element wrapper. Note that after the keyword **existing** or after the driven element type, a colon is used as a separator. This is probably the only exception to the rule that the colon can only follow a **description** at the beginning of a card. The label **driven_elem_label** of the driven element must match that of the driving element used at the beginning of the **driven** card.

Example. A pin constraint between two rigid bodies is released:

```
set: integer BODY_1 = 1;
set: integer BODY_2 = 2;
# ...
structural: BODY_1, dynamic,
  null, eye, null, null;
structural: BODY_2, dynamic,
  null, eye, null, null;
# ....
body: BODY_1, BODY_1,
  1000., null, diag, 100.,100.,1.;
body: BODY_2, BODY_2,
  10., null, diag, 1.e-1,1.e-1,1.e-3;
# ...
# this constraint will be released when Time = 10 s
driven: 1, string, "Time < 10.",
joint: 1, spherical hinge,
  BODY_1,
  position, null,
  BODY_2,
  position, null;
```

Example. an **axial rotation** (see Section 8.12.3) joint is replaced by a **revolute hinge** (see Section 8.12.38) when the desired spin velocity, measured as the angular velocity of the second node (assuming, for instance, that the first one is fixed), is reached. The value of an abstract node is used to input the angular velocity to the **axial rotation** joint.

```
set: integer BODY_1 = 1;
set: integer BODY_2 = 2;
set: integer CONTROL_OUTPUT = 3;
# ...
structural: BODY_1, static,
  null, eye, null, null;
structural: BODY_2, dynamic,
  null, eye, null, null;
abstract: CONTROL_OUTPUT;
# ....
driven: 1, node, BODY_2, structural, string, "Omega[3]",
  string, "Var < 100.",
joint: 1, axial rotation,
  BODY_1,
```

```

        position, null,
        orientation, 1, 1.,0.,0., 3, 0.,0.,1.,
BODY_2,
        position, null,
        orientation, 1, 1.,0.,0., 3, 0.,0.,1.,
node, CONTROL_OUTPUT, abstract, string, "x",
        linear, 0.,1.;
driven: 2, node, BODY_2, structural, string, "Omega[3]",
        string, "Var >= 100.",
joint: 2, revolute hinge,
BODY_1,
        position, null,
        orientation, 1, 1.,0.,0., 3, 0.,0.,1.,
BODY_2,
        position, null,
        orientation, 1, 1.,0.,0., 3, 0.,0.,1.;

```

Hint

The **hint** feature consists in allowing the setup of elements to be computed when they are activated rather than at startup. For example, a joint that is activated after some simulation time may need to compute its relative position and orientation from the parameters of the simulation; a drive that controls the evolution of the relative configuration of a joint may need to infer its parameters from the current configuration of the overall system; and so on.

Currently, only few elements, significantly joints, support and honor hints. The typical syntax is given by a keyword followed by some optional parameters within curly brackets. For instance, the hint that instructs a joint to compute the offset with respect to the second node when it is activated is

```

driven: 10, string, "Time > 1.", hint, "offset{2}",
joint: 10, ...

```

A similar form is used to instruct a joint to compute the relative orientation with respect to node 1:

```

driven: 20, string, "Time > 1.", hint, "hinge{1}",
joint: 20, ...

```

A **distance** joint may be fed a new drive by using:

```

set: const real T0 = 1.;
driven: 30, string, "Time > T0",
        hint, "drive{cosine, T0, pi/.1, -.9/2., half, model::distance(1,2)}",
joint: 30, distance, ...

```

A **drive hinge** joint may be fed a new drive by using:

```

set: const real T0 = 1.;
driven: 40, string, "Time > T0",
        # note: the lines wrap for typographical reasons
        # in the actual file, the string has to be a single line
        hint, "drive3{model::xdistance(1,2), \
                model::ydistance(1,2), model::zdistance(1,2), \
                cosine, T0, pi/.1, -.9/2., half, model::distance(1,2)}",
joint: 30, drive hinge, ...

```

This feature will likely be extended to other elements and generalized as much as possible.

8.17.3 Inertia

This card causes the equivalent inertia properties of a subset of elements to be generated.

```
<card> ::= inertia : <label>
    [ , name , " <inertia_name> " ]
    [ , position , (Vec3) <reference_position> ]
    [ , orientation , (Mat3x3) <reference_orientation> ]
    , <type_subset> [ , ... ]
    [ , output , { no | yes | log | both } ] ;

<type_subset> ::= <type> , { all | <label> [ , ... ] }

<type> ::= { body | joint | loadable }
```

where `type` currently can be `body`, `joint` and `loadable`, although more elements associated to inertia might participate in the future. All elements whose labels are listed must exist, and duplicates are detected and considered errors. The keyword `all` causes all the elements of type `type` to be included in the list.

The only effect of the `inertia` statement is to log each `inertia` entry in the `.log` file in a self explanatory form, referred both to the `global` reference frame and to a reference frame originating from `reference_position` and oriented as `reference_orientation`. The optional parameter `output` may be used to alter the default behavior:

- `no` disables the output, making the `inertia` statement essentially ineffective;
- `yes` enables output to standard output;
- `log` enables output to the `.log` file (the default);
- `both` enables output to both standard output and `.log` file.

8.17.4 Output

This card does not instantiate any `element`; it rather enables output of selected elements, and it is analogous to that of the `nodes` (see Section 6.6.1):

```
<card> ::= output : <elem_type> , <elem_list> ;

<elem_list> ::=
    { <elem_label> [ , ... ]
      | range , <elem_start_label> , <elem_end_label> }
```

`elem_type` is a valid element type that can be read as card name in the `elements` block. In case the keyword `range` is used, all existing elements comprised between `elem_start_label` and `elem_end_label` are set; missing ones are silently ignored.

Appendix A

NetCDF Output Format

NetCDF support has been initially contributed by Patrick Rix.

NetCDF is a format to efficiently store and retrieve data to and from a database on file in a portable, platform-independent manner. Further details can be found at

<http://www.unidata.ucar.edu/software/netcdf/>

and in [18]. The link reported above also describes some of the common tools that can be used to read and manipulate the contents of the database.

The output in NetCDF format consists in a single binary file written by the NetCDF library and intended to be read by tools exploiting the library itself. This document does not describe the details of NetCDF low-level format, since this is not intended to be accessed directly by MBDyn users. Interested readers can consult the specific documentation [18].

The output in NetCDF format is a work-in-progress, so it may be subjected to changes during the development of MBDyn. This chapter is intended to document its current status, so it may be incomplete, and occasionally outdated as well.

A.1 NetCDF Output

Following the convention of NetCDF data, each datum is defined by a variable, whose name indicates the type of datum and the entity that generated it, organized in a tree-like fashion.

For example, the vector containing the three components of the position of the **structural** node labeled **label** is

```
node.struct.<label>.X
```

Each variable usually has few attributes:

- a **description** is usually given;
- the **units** are specified, unless variables are non-dimensional;
- the **type** is specified, if relevant; it contains the name of the C/C++ structure the data was taken from.

Each level contains some datum that is intended to make the contents of the database as self-explanatory as possible. For example, the level **node** contains an array whose values are the strings indicating the node types available; each level **node.<type>** contains the labels of the available nodes for that type, and so on.

A.1.1 Base Level

Currently, the following basic levels are available:

- `run`, for general simulation-related data;
- `node`, for nodes;
- `elem`, for elements.

A.1.2 Run Level

There is no variable `run`, with the name of the run level. This level contains general, simulation-related data:

- `run.step`, the step number;
- `time` (previously `run.time`), the value of the time;
- `run.timestep`, the value of the timestep.

Note: the `run.time` variable has been renamed `time` because many NetCDF manipulation tools automatically recognize this name as the ordering variable for the rest of the data.

A.1.3 Node Level

There is no variable `node`, with the name of the node level. This level contains as many sublevels as the node types that are present in the model. Each node type consists in a variable that contains an array with the labels (integers) of the nodes defined for that type. A variable named `node.<type>.<label>` may be present (it is, for example, for `structural nodes`); however, useful data is usually contained in specific variables, whose names describe the contents.

Currently supported node types are:

- Structural nodes, indicated as `struct`.

Structural Node

The following variables are defined.

- `node.struct.<label>` is actually empty; it contains the type of the structural node in the `type` attribute;
- `node.struct.<label>.X` contains the position of the node, in the appropriate reference frame;
- `node.struct.<label>.R` contains the orientation matrix of the node, in the appropriate reference frame;
- `node.struct.<label>.Phi` contains the orientation vector describing the orientation of the node, in the appropriate reference frame;
- `node.struct.<label>.E` contains the Euler angles describing the orientation of the node, in the appropriate reference frame;
- `node.struct.<label>.XP` contains the velocity of the node, in the appropriate reference frame;
- `node.struct.<label>.Omega` contains the angular velocity of the node, in the appropriate reference frame.

Note that only one of **R**, **Phi**, or **E** are present, depending on the requested description of the orientation of the node (see **structural node**, Section 6.1, and **default orientation**, Section 5.2.10).

The **dynamic** and **modal** node types, if requested, can output extra variables.

- **node.struct.<label>.XPP** contains the acceleration of the node, in the appropriate reference frame;
- **node.struct.<label>.OmegaP** contains the angular acceleration of the node, in the appropriate reference frame.

If requested, the inertia associated to **dynamic** nodes is output within the namespace of the node, although actually handled by the corresponding **automatic structural** element. As a consequence, extra variables can appear in the output of **dynamic** nodes.

- **node.struct.<label>.B** contains the momentum of the inertia associated to the node, in the appropriate reference frame;
- **node.struct.<label>.G** contains the momenta moment of the inertia associated to the node, in the appropriate reference frame;
- **node.struct.<label>.BP** contains the derivative of the momentum of the inertia associated to the node, in the appropriate reference frame;
- **node.struct.<label>.GP** contains the derivative of momenta moment of the inertia associated to the node, in the appropriate reference frame.

A.1.4 Element Level

There is no variable **elem**, with the name of the element level. This level contains as many sublevels as the element types that are present in the model. Each element type consists in a variable that contains an array with the labels (integers) of the elements defined for that type. A variable named **elem.<type>.<label>** may be present; however, useful data is usually contained in specific variables, whose names describe the contents.

Currently supported element types are:

- Automatic structural elements, indicated as **autostruct**.
- Beam elements, indicated as **beam2** and **beam3** for 2- and 3-node beam elements, respectively.

Aerodynamic Elements

NetCDF output of aerodynamic elements has been sponsored by REpower Systems AG.

The aerodynamic elements allow to define a broad set of variables. Values are output at the integration points. Each element allows to define n integration points. The **aerodynamic body** outputs at n points. The **aerodynamic beam2** outputs at $2 \cdot n$ points. The **aerodynamic beam3** outputs at $3 \cdot n$ points.

Output.

- **elem.aerodynamic.<label>.X_<i>** contains the location of the integration point **i**, in the global reference frame;
- **elem.aerodynamic.<label>.R_<i>** contains the orientation matrix of the airfoil at the integration point **i**, in the global reference frame;

- `elem.aerodynamic.<label>.Phi_<i>` contains the rotation vector that describes the orientation of the airfoil at the integration point `i`, in the global reference frame;
- `elem.aerodynamic.<label>.E_<i>` contains the set of Euler angles that describes the orientation of the airfoil at the integration point `i`, in the global reference frame. The sequence of the angles is detailed in the variable's description. It can be 123 (the default), 313 or 321;
- `elem.aerodynamic.<label>.V_<i>` contains the velocity of the integration point `i`, in the global reference frame;
- `elem.aerodynamic.<label>.Omega_<i>` contains the angular velocity of the integration point `i`, in the global reference frame;
- `elem.aerodynamic.<label>.F_<i>` contains the force at the integration point `i`, in the absolute reference frame;
- `elem.aerodynamic.<label>.M_<i>` contains the moment at the integration point `i`, in the absolute reference frame, referred to the location of the integration point.

Note: variables `R_<i>`, `Phi_<i>` and `E_<i>` are mutually exclusive.

Note: by default, only variables `elem.aerodynamic.<label>.F_<i>` and `elem.aerodynamic.<label>.M_<i>` are output. See Section 8.1.1 for indications about how to customize aerodynamic element output in NetCDF format.

Automatic Structural

No variables are explicitly defined for the automatic structural element; on the contrary, their specific data, if requested, is appended to the corresponding dynamic structural node.

Beam

The beam elements allow to define a broad set of variables.

Two-Node Beam Element. The two-node beam element allows to define the variables

- `elem.beam.<label>.X` contains the location of the evaluation point, in the global reference frame;
- `elem.beam.<label>.R` contains the orientation matrix of the beam section at the evaluation point, in the global reference frame;
- `elem.beam.<label>.Phi` contains the rotation vector that describes the orientation of the beam section at the evaluation point, in the global reference frame;
- `elem.beam.<label>.E` contains the set of Euler angles that describes the orientation of the beam section at the evaluation point, in the global reference frame. The sequence of the angles is detailed in the variable's description. It can be 123 (the default), 313 or 321;
- `elem.beam.<label>.F` contains the internal force at the evaluation point, in the reference frame of the beam section;
- `elem.beam.<label>.M` contains the internal moment at the evaluation point, in the reference frame of the beam section;
- `elem.beam.<label>.mu` contains the linear strain at the evaluation point, in the reference frame of the beam section;

- `elem.beam.<label>.k` contains the angular strain at the evaluation point, in the reference frame of the beam section;
- `elem.beam.<label>.muP` contains the linear strain rate at the evaluation point, in the reference frame of the beam section;
- `elem.beam.<label>.kP` contains the angular strain rate at the evaluation point, in the reference frame of the beam section.

Note: variables `R`, `Phi` and `E` are mutually exclusive.

Note: variables `muP` and `kP` are only available for viscoelastic beam elements.

Note: by default, only variables `F` and `M` are output. See Section 8.3.2 for indications about how to customize beam element output in NetCDF format.

Three-Node Beam Element. The three-node beam element allows to define the same set of variables of the two-node one, postfixed with either `_I` or `_II` to indicate the first (between nodes 1 and 2) or the second (between nodes 2 and 3) evaluation point.

A.2 Accessing the Database

The database can be accessed using any of the tools listed at the web site

<http://www.unidata.ucar.edu/software/netcdf/software.html>

(yes, including MBDyn itself...)

A.2.1 Octave

The `octcdf` package, available from

<http://ocgmod1.marine.usf.edu/>,

provides a clean interface to using NetCDF databases from within the popular math environment Octave. Results from MBDyn can be easily handled once the data structure is known.

To access the database, a handler needs to be obtained by calling

```
octave:1> nc = netcdf('mbdyn_output.nc', 'r');
```

Variable descriptions are accessed as

```
octave:2> nc{'node.struct.1000.X'}
```

Their values are accessed as

```
octave:3> nc{'node.struct.1000.X'}(10, 3)
```

So, for example, the z component of the position of node 1000 can be plot with

```
octave:4> plot(nc{'time'}(:), nc{'node.struct.1000.X'}(:,3))
```

Appendix B

Modal Element Data

B.1 FEM File Format

This section describes the format of the FEM input to the `modal` joint of MBDyn. This data can be obtained, for example:

- from Code Aster (<http://www.code-aster.org/>), as discussed in Section B.2;
- from NASTRAN (<http://www.mscsoftware.com/>) output, using the `femgen` utility, as detailed in Appendix B.3 (in short, it processes binary output from NASTRAN, as defined by means of appropriate ALTER files provided with MBDyn sources, into a file that is suitable for direct input in MBDyn);
- by manually crafting the output of your favourite FEM analysis: since it is essentially a plain ASCII file, it can be generated in a straightforward manner from analogous results obtained with almost any FEM software, from experiments or manually generated from analytical or numerical models of any kind.

The format is:

```
<comments>
** RECORD GROUP 1,<any comment to EOL; "HEADER">
<comment>
<REV> <NNODES> <NNORMAL> <NATTACHED> <NCONSTRAINT> <NREJECTED>
<comments; NNODES = NNORMAL + NATTACHED + NCONSTRAINT - NREJECTED>
** RECORD GROUP 2,<any comment to EOL; "FINITE ELEMENT NODE LIST">
<FEMLABEL> [...NNODES]
<comments; FEM label list: NNODES integers>
** RECORD GROUP 3,<any comment to EOL; "INITIAL MODAL DISPLACEMENTS">
<MODEDISP> [...NNODES]
<comments; initial mode displacements: NNODES reals>
** RECORD GROUP 4,<any comment to EOL; "INITIAL MODAL VELOCITIES">
<MODEVEL> [...NNODES]
<comments; initial mode velocities: NNODES reals>
** RECORD GROUP 5,<any comment to EOL; "NODAL X COORDINATES">
<FEM_X> [...NNODES]
<comments; FEM node X coordinates>
** RECORD GROUP 6,<any comment to EOL; "NODAL Y COORDINATES">
<FEM_Y> [...NNODES]
<comments; FEM node Y coordinates>
```

```

** RECORD GROUP 7,<any comment to EOL; "NODAL Z COORDINATES">
<FEM_Z> [...NNODES]
<comments; FEM node Z coordinates>
** RECORD GROUP 8,<any comment to EOL; "NON-ORTHOGONALIZED MODE SHAPES">
<comment; NORMAL MODE SHAPE # 1>
  <FEM1_X> <FEM1_Y> <FEM1_Z> <FEM1_RX> <FEM1_RY> <FEM1_RZ>
  [...NNODES]
<comment; NORMAL MODE SHAPE # 2>
  <FEM2_X> <FEM2_Y> <FEM2_Z> <FEM2_RX> <FEM2_RY> <FEM2_RZ>
  [...NNODES]
[...NMODES]
<comments; for each MODE, for each NODE: modal displacements/rotations>
** RECORD GROUP 9,<any comment to EOL; "MODAL MASS MATRIX. COLUMN-MAJOR FORM">
<M_1_1>      [...] <M_1_NMODES>
[...]
<M_NMODES_1> [...] <M_NMODES_NMODES>
<comments; the modal mass matrix in column-major (symmetric?)>
** RECORD GROUP 10,<any comment to EOL; "MODAL STIFFNESS MATRIX. COLUMN-MAJOR FORM">
<K_1_1>      [...] <K_1_NMODES>
[...]
<K_NMODES_1> [...] <K_NMODES_NMODES>
<comments; the modal stiffness matrix in column-major (symmetric?)>
** RECORD GROUP 11,<any comment to EOL; "DIAGONAL OF LUMPED MASS MATRIX">
<M_1_X> <M_1_Y> <M_1_Z> <M_1_RX> <M_1_RY> <M_1_RZ>
[...]
<M_NNODES_X> [...] <M_NNODES_RZ>
<comments; the lumped diagonal mass matrix of the FEM model>
** RECORD GROUP 12,<any comment to EOL; "GLOBAL INERTIA PROPERTIES">
<MASS>
<X_CG> <Y_CG> <Z_CG>
<J_XX> <J_XY> <J_XZ>
<J_YX> <J_YY> <J_YZ>
<J_ZX> <J_ZY> <J_ZZ>
<comments; the global inertia properties of the modal element>

```

An arbitrary number of comment lines may appear where <comments[...]> is used; only one comment line must appear where where <comment[...]> is used.

The beginning of a record is marked

```

** RECORD GROUP <RID>

```

where the number **RID** indicates what record is being read. The size of each record, i.e. the number of values that are expected, is defined based on the header record, so MBDyn should be able to detect incomplete or mis-formatted files.

The records contain:

- **RECORD GROUP 1**, a.k.a. the “header”, contains a summary of the contents of the file:
 - **REV** is a string that indicates the revision number; it is currently ignored;
 - **NNODES** is the number of (exposed) FEM nodes in the FEM model;
 - **NNORMAL** is the number of normal modes;
 - **NATTACHED** is the number of “attached”, i.e. static, modes;
 - **NCONSTRAINT** is the number of constraint modes;

- `NREJECTED` is the number of rejected modes.

Currently, the number of available modes is computed as

$$\text{NMODES} = \text{NNORMAL} + \text{NATTACHED} + \text{NCONSTRAINT} - \text{NREJECTED}$$

because modes are treated in a generalized manner, so there is no need to consider the different types of shapes in a specific manner. Typically, one should set all numbers to 0, except for `NNORMAL` which should be set to the number of modes actually present in the data set. Remember that MBDyn can still select a subset of the available modes to be used in the analysis, so that there is no need to further edit this file.

- **RECORD GROUP 2** contains a listing of the `NNODES` labels of the (exposed) FEM nodes. The labels can be any string of text, and are separated by blanks (as intended by the `isspace(3)` C standard library function). A label cannot thus contain any amount of whitespace.
- **RECORD GROUP 3** contains the initial values of the `NMODES` modal unknowns (optional; set to zero if omitted);
- **RECORD GROUP 4** contains the initial values of the `NMODES` modal unknowns derivatives (optional; set to zero if omitted);
- **RECORD GROUP 5** contains the X component of the position of the `NNODES` FEM nodes in the reference frame attached to the `modal` node (or to the `origin node`, if given).
- **RECORD GROUP 6** contains the Y component of the data above;
- **RECORD GROUP 7** contains the Z component of the data above;
- **RECORD GROUP 8** contains the non-orthogonalized components of the `NMODES` modes; for each mode, the three components X , Y , Z of the modal displacement, and the three components RX , RY , RZ of the linearized modal rotations are listed; each mode shape is separated by a comment line, which typically is

```
**    NORMAL MODE SHAPE #    <N>
```

for readability;

- **RECORD GROUP 9** contains the modal mass matrix, i.e. a square, `NMODES` × `NMODES` matrix that contains the reduced mass $[m]$ resulting from the multiplication

$$[m] = \{X\}^T [M] \{X\} \tag{B.1}$$

When only normal modes are used, it is diagonal. It can be semi-definite positive, or even zero, if a partially or entirely static model is considered.

- **RECORD GROUP 10** contains the modal stiffness matrix, i.e. a square, `NMODES` × `NMODES` matrix that contains the reduced stiffness $[k]$ resulting from the multiplication

$$[k] = \{X\}^T [K] \{X\} \tag{B.2}$$

When only normal modes are used, it is diagonal; in that case, the diagonal contains the modal mass times the square of the eigenvalues, i.e. $k_{ii} = \omega_i^2 m_{ii}$. It should be definite positive; in fact, rigid degrees of freedom that would make it semi-definite should rather be modeled by combining separate modal submodels by way of multibody connections, so that the multibody capability to handle finite relative displacements and rotations is exploited. Note however that the positive-definiteness of the generalized stiffness matrix is not a requirement.

- **RECORD GROUP 11** contains the lumped inertia matrix associated to the **NNODES** (exposed) FEM nodes; for each node, the X , Y , Z , RX , RY and RZ inertia coefficients are listed (they can be zero). The nodal mass coefficients Y and Z should be equal to coefficient X . The resulting diagonal matrix should satisfy the constraint illustrated in Equation B.1.
- **RECORD GROUP 12** contains the global inertia properties of the modal element: the total mass, the three components of the position of the center of mass in the local frame of the FEM model, and the 3×3 inertia matrix with respect to the center of mass, which is supposed to be symmetric and positive definite (or semi-definite, if the model is static).

Note that **RECORD GROUP 11** and **RECORD GROUP 12** used to be mutually exclusive in earlier versions. The reason for accepting **RECORD GROUP 12** format, regardless of **RECORD GROUP 11** presence, is related to the fact that in many applications the FEM nodal inertia may not available and, at the same time, a zero-order approximation of the inertia is acceptable. The reason for allowing both is that when **RECORD GROUP 11** is present, its data used to be used to compute the rigid-body motion invariants that can be provided by **RECORD GROUP 12**. This might represent an unacceptable approximation. **RECORD GROUP 12** allows to replace those invariants by a better estimate, when available.

Both **RECORD GROUP 11** and **RECORD GROUP 12** may be absent. This case only makes sense when a zero-order approximation of the inertia is acceptable and the rigid-body motion of the modal element is not allowed, i.e. the element is clamped to the ground (see the keyword **clamped** in Section 8.12.32 for details).

Although the format loosely requires that no more than 6 numerical values appear on a single line, MBDyn is very forgiving about this and can parse the input regardless of the formatting within each record. However this liberality may lead to inconsistent or unexpected parsing behavior, so be careful.

B.1.1 Example: Dynamic Model

As an example, a very simple, hand-made FEM model file is presented below. It models a FEM model made of three aligned nodes, where inertia is evenly distributed. Note that each line is prefixed with a two-digit line number that is not part of the input file. Also, for readability, all comments are prefixed by “**”, in analogy with the mandatory “** RECORD GROUP” lines, although not strictly required by the format of the file.

```

01  ** MBDyn MODAL DATA FILE
02  ** NODE SET "ALL"
03
04
05  ** RECORD GROUP 1, HEADER
06  **  REVISION,  NODE,  NORMAL,  ATTACHMENT,  CONSTRAINT,  REJECTED MODES.
07      REVO      3      2      0      0      0
08  **
09  ** RECORD GROUP 2, FINITE ELEMENT NODE LIST
10      1001 1002 1003
11
12  **
13  ** RECORD GROUP 3, INITIAL MODAL DISPLACEMENTS
14      0 0
15  **
16  ** RECORD GROUP 4, INITIAL MODALVELOCITIES
17      0 0
18  **
19  ** RECORD GROUP 5, NODAL X COORDINATES

```

```

20  0
21  0
22  0
23  **
24  ** RECORD GROUP 6, NODAL Y COORDINATES
25  -2.
26  0
27  2.
28  **
29  ** RECORD GROUP 7, NODAL Z COORDINATES
30  0
31  0
32  0
33  **
34  ** RECORD GROUP 8, MODE SHAPES
35  **   NORMAL MODE SHAPE #  1
36  0 0 1 0 0 0
37  0 0 0 0 0 0
38  0 0 1 0 0 0
39  **   NORMAL MODE SHAPE #  2
40  1 0 0 0 0 0
41  0 0 0 0 0 0
42  1 0 0 0 0 0
43  **
44  ** RECORD GROUP 9, MODAL MASS MATRIX
45  2 0
46  0 2
47  **
48  ** RECORD GROUP 10, MODAL STIFFNESS MATRIX
49  1 0
50  0 1e2
51  **
52  ** RECORD GROUP 11, DIAGONAL OF LUMPED MASS MATRIX
53  1 1 1 1 1 1
54  1 1 1 1 1 1
55  1 1 1 1 1 1

```

The corresponding global inertia properties are:

$$m = 3 \tag{B.3}$$

$$\mathbf{x}_{CM} = \begin{Bmatrix} 0 \\ 0 \\ 0 \end{Bmatrix} \tag{B.4}$$

$$\mathbf{J} = \begin{bmatrix} 11 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 11 \end{bmatrix} \tag{B.5}$$

B.1.2 Example: Static Model

As an example, a very simple, hand-made FEM model file is presented below. It models a FEM model made of three aligned nodes, where inertia is only associated to the mid-node. As a consequence, the three mode shapes must be interpreted as static shapes, since the modal mass matrix is null. Note that each line is prefixed with a two-digit line number that is not part of the input file. Also, for readability,

all comments are prefixed by “**”, in analogy with the mandatory “** RECORD GROUP” lines, although not strictly required by the format of the file.

```
01 ** MBDyn MODAL DATA FILE
02 ** NODE SET "ALL"
03
04
05 ** RECORD GROUP 1, HEADER
06 ** REVISION, NODE, NORMAL, ATTACHMENT, CONSTRAINT, REJECTED MODES.
07 REVO      3      2      0      0      0
08 **
09 ** RECORD GROUP 2, FINITE ELEMENT NODE LIST
10 1001 1002 1003
11
12 **
13 ** RECORD GROUP 3, INITIAL MODAL DISPLACEMENTS
14 0 0
15 **
16 ** RECORD GROUP 4, INITIAL MODALVELOCITIES
17 0 0
18 **
19 ** RECORD GROUP 5, NODAL X COORDINATES
20 0
21 0
22 0
23 **
24 ** RECORD GROUP 6, NODAL Y COORDINATES
25 -2.
26 0
27 2.
28 **
29 ** RECORD GROUP 7, NODAL Z COORDINATES
30 0
31 0
32 0
33 **
34 ** RECORD GROUP 8, MODE SHAPES
35 ** NORMAL MODE SHAPE # 1
36 0 0 1 0 0 0
37 0 0 0 0 0 0
38 0 0 1 0 0 0
39 ** NORMAL MODE SHAPE # 2
40 1 0 0 0 0 0
41 0 0 0 0 0 0
42 1 0 0 0 0 0
43 **
44 ** RECORD GROUP 9, MODAL MASS MATRIX
45 0 0
46 0 0
47 **
48 ** RECORD GROUP 10, MODAL STIFFNESS MATRIX
49 1 0
50 0 1e2
51 **
```

```

52 ** RECORD GROUP 12, GLOBAL INERTIA
53 3
54 0 0 0
55 11 0 0
56 0 3 0
57 0 0 11

```

The same global inertia properties of the model in Section B.1.1 have been used; as a result, the two models show the same rigid body dynamics behavior, but the dynamic model also shows deformable body dynamic behavior, while the static one only behaves statically when straining is involved.

B.2 Code Aster Procedure

The `fem_data_file` can be generated by Code Aster, using the macro provided in the `cms.py` file, which is located in directory `contrib/CodeAster/cms/` of the distribution, and is installed in directory `$PREFIX/share/py/`.

Preparing the input for Code Aster essentially requires to prepare the bulk of the mesh either manually or using some meshing tools (e.g. `gmsh`), as illustrated in the related documentation and tutorials, and then writing a simple script in `python` (<http://www.python.org/>) that defines the execution procedure.

To produce the data for MBDyn's modal element, a specific macro needs to be invoked during the execution of Aster. The macro instructs Aster about how the solution procedure needs to be tailored to produce the desired results, and then generates the `.fem` file with all the requested data.

The steps of the procedure are as follows:

1. prepare a Code Aster input model (the `.mail` file), containing the nodes, the connections and at least a few groups of nodes, as indicated in the following;
2. prepare a Code Aster command file (the `.comm` file), following, for example, one of the given templates;
3. as soon as the model, the materials, the element properties and any optional boundary constrains are defined, call the `CMS` macro with the appropriate parameters. The syntax of the macro is:

```

CMS(
  MAILLAGE = <mesh_concept> ,
  [ INTERFACE = <interface_node_set_name> , ]
  EXPOSED = _F( { GROUP_NO = <exposed_node_set_name> | TOUT = 'OUI' } ),
  MODELE = <model_concept> ,
  CARA_ELEM = <element_properties_concept> ,
  CHAM_MATER = <material_properties_concept> ,
  [ CHAR_MECA = <boundary_conditions_concept> ,
  [ OPTIONS = _F(
    [ NMAX_FREQ = <maximum_number_of_frequencies> ]
  ), ]
  OUT = _F(
    TYPE = { 'MBDYN' } ,
    [ PRECISION = <number_of_digits> , ]
    [ DIAG_MASS = { 'OUI' | 'NON' } , ]
    [ RIGB_MASS = { 'OUI' | 'NON' } , ]
    FICHIER = <output_file_name>
  )
);

```

where:

- **MAILLAGE** is the name of a python object returned by **LIRE_MAILLAGE**;
- **INTERFACE** is a string containing the name of the group of nodes, defined in the mesh file, that will be used as interface to add Craig-Bampton shapes to the base (if omitted, only normal modes are used);
- **EXPOSED** allows two formats:
 - **GROUP_NO** is a string containing the name of the group of nodes, defined in the mesh file, that will be exposed, i.e. that will appear in the `.fem` file generated by this macro;
 - **TOUT** is a string that must be set to `'OUI'`, indicating that all nodes will be exposed (technically speaking, all nodes involved in at least one connection);
- **MODELE** is the name of a python object returned by **AFFE_MODELE**;
- **CARA_ELEM** is the name of a python object returned by **AFFE_CARA_ELEM**;
- **CHAM_MATER** is the name of a python object returned by **AFFE_MATERIAU**;
- **CHAR_MECA** is the name of a python object returned by **AFFE_CHAR_MECA** (optional);
- **OPTIONS** allows a set of generic options to be passed:
 - **NMAX_FREQ** is the number of frequencies expected from the eigenanalysis, and thus the number of normal modes that will be added to the modal base; the smallest frequencies are used;
- **OUT** describes the parameters related to the type of output the macro generates; right now, only output in the format specified in this section is available:
 - **TYPE** indicates the type of output the macro is requested to generate; it must be `'MBDYN'`;
 - **FICHIER** (“file” in French) is the name of the file where the generated data must be written; it must be a fully qualified path, otherwise it is not clear where the file is actually generated;
 - **PRECISION** contains the number of digits to be used when writing the data (for example, **PRECISION** = 8 results in using the format specifier `%16.8e`);
 - **DIAG_MASS** requests the output of the diagonal of the mass matrix (RECORD GROUP 11); this option really makes sense only if all nodes are output (namely, if the nodes set passed to **EXPOSED** consists in all nodes);
 - **DIAG_MASS** requests the output of the rigid-body inertia properties (RECORD GROUP 12);

the macro returns a python object consisting in the generalized model it used to generate the file, as returned by **MACR_ELEM_DYNA**.

4. prepare a `.export` file, either following the examples or using `astk`;
5. run Aster.

Annotated examples are provided in directory `contrib/CodeAster/cms/`.

B.3 NASTRAN Procedure

The `fem_data_file` can be generated using NASTRAN, by means of the **ALTER** cards provided in directory `etc/modal.d/` of the distribution.

The steps of the procedure are as follows:

1. prepare a NASTRAN input card deck, made of bulk data, eigenanalysis properties data and/or static loads (details about this phase are currently missing and will be provided in future releases of the manual).
2. complete the NASTRAN input file by putting some specific **ALTER** cards. In detail:
 - (a) the file `MBDyn_NASTRAN_alter_1.nas` contains **ALTER** definitions for static solutions; appropriate loading subcases for each solution must be provided in the case control and in the bulk data sections of the input file;
FIXME: I don't know how to use static shapes only.
 - (b) the file `MBDyn_NASTRAN_alter_2.nas` contains **ALTER** definitions for eigenanalysis solutions; an appropriate eigenanalysis method, with the related data card must be provided in the case control and in the bulk data sections of the input file;
FIXME: I don't know how to use this together with static shapes; I only get the normal mode shapes, even if the matrices are complete.
 - (c) the file `MBDyn_NASTRAN_alter_3.nas` contains **ALTER** definitions for eigenanalysis solutions; an appropriate eigenanalysis method, with the related data card must be provided in the case control and in the bulk data sections of the input file;
Note: this works; see tests/modal/beam.README.

Exactly one of these files must be included at the very top of the NASTRAN input file; they already include the appropriate **SOL** statement, so the input file must begin with

```
$ Replace '#' below with number that matches your needs
INCLUDE 'MBDyn_NASTRAN_alter_#.nas'
CEND
$... any other executive control and bulk data card
```

The static solution of case (a: **SOL 101**) and the eigensolution of case (b: **SOL 103**) need to be performed in sequence; if only the eigensolution is to be used, the **ALTER** file of case (c: **SOL 103**) must be used. The static solution of case (a) generates a binary file `mbdyn.stm`; the eigensolutions of cases (b-c) generate two binary files, `mbdyn.mat` and `mbdyn.tab`, which, in case (b), include the static solutions as well. The **ALTER** currently included in the MBDyn distribution work correctly only with the following **PARAM** data card:

```
PARAM,POST,-1
```

3. Run NASTRAN.
4. Run the tool `femgen`, which transforms the above binary files into the `fem_data_file`. The file name is currently requested as terminal input; this is the name of the file that will be used in the input model for `mbdyn`. Conventionally, the `.fem` extension is used.

B.4 Procedures for Other FEA Software

Currently, no (semi-)automatic procedures are available to extract modal element data from other FEA software. If you need to use any other software, and possibly a free software, you'll need to design your own procedure, which may go from a very simple script (shell, awk, python) for handling of textual output, to a fully featured modeling and translation interface. If you want to share what you developed with the MBDyn project, feel free to submit code, procedures or just ideas and discuss them on the `mbdyn-users@mbdyn.org` mailing list.

Appendix C

Results Visualization

This section describes the different ways the raw output from MBDyn can be arranged to work with third-party software for visualization. In most of the cases, the preparation of the results can be done as a post-processing, starting from raw MBDyn output files. This is true, for instance, for EasyAnim (see <http://mecara.fpms.ac.be/EasyDyn/> for details).

In previous versions of MBDyn, the `output results` statement allowed to generate output compatible with MSC.ADAMS; provisions for Altair Motion View has been under preparation for long time, but it has never been truly supported. Anyone interested in that interface should contact MBDyn developers. The `output results` statement is now deprecated; support for all visualization tools will be reworked in form of postprocessing of MBDyn's raw output, either in textual or binary form. Right now, that statement is used to enable the experimental support for NetCDF.

Some post-processing preparation instructions are available for those packages that require special handling and thus are built-in.

C.1 EasyAnim

MBDyn exploits a special version of EasyAnim, based on 1.3 sources and patched by MBDyn developers; the patch is known to work under UN*X (GNU/Linux, essentially); it has not been tested with Windows, except with `cygwin`. It is available from the MBDyn web site, and it has been submitted to EasyAnim developers for evaluation. Check out MBDyn's web site for more information.

The preparation of the output is done via `awk(1)`, which is invoked by the `mbdyn2easyanim.sh` script. It uses the `.log` and `.mov` files, and generates a pair of `.vol` and `.van` files which contain the model and the motion, respectively. Use

```
mbdyn2easyanim.sh <mbdyn_output_filename>
```

where `mbdyn_output_filename` is the name of the output file from MBDyn, with no extension. This script generates a model based on the topology information contained in the `mbdyn_output_filename.log` file, plus the animation information according to the contents of the `mbdyn_output_filename.mov` file.

To check the model during preparation, run MBDyn with

```
abort after: input;
```

in the `problem_name` section (pag 95), so that the model is output in the input configuration, and EasyAnim can visualize it.

The model can be customized by supplying the `mbdyn2easyanim.sh` script additional `awk(1)` code that adds nodes, edges and sides whose movement is associated to that of existing nodes.

For this purpose, one needs to write an `awk(1)` script like

```

BEGIN {
    # add a node property
    # node prop name, color, radius
    nodeprop_add("dummy_node", 4, 1.);

    # add a node
    # node name, base node name, x, y, z in base node frame, prop name
    node_add("added_node", "base_node", 1., 2., 3., "dummy_node");

    # add an edge property
    # edge prop name, color, radius
    edgeprop_add("dummy_edge", 15, 2.);

    # add an edge
    # edge name, node 1 name, node 2 name, prop name
    edge_add("added_edge", "node_1", "node_2", "dummy_edge");

    # add a side property
    # side prop name, color
    sideprop_add("dummy_side", 8);

    # add a side
    nodes[1] = "node_1";
    nodes[2] = "node_2";
    nodes[3] = "node_3";
    nodes[4] = "node_4";
    # side name, number of nodes, node labels, prop name
    side_add("added_side", 4, nodes, "dummy_side");
}

```

place it into a file, say `custom.awk`, and then invoke the interface script as

```
mbdyn2easyanim.sh -f custom.awk <mbdyn_output_filename>
```

C.2 Output Manipulation

This section notes how output can be manipulated in useful manners.

C.2.1 Output in a Relative Reference Frame

In many cases it may be useful to represent the output of **structural** nodes in a reference frame relative to one node. For example, to investigate the deformation of a helicopter rotor, it may be useful to refer the motion of the blades to the reference frame of the rotor hub, so that the reference rotation of the rotor is separated from the local straining.

For this purpose, MBDyn provides an `awk(1)` script that extracts the desired information from the `.mov` file.

The script is `$(SRC)/var/abs2rel.awk` and is usually installed in `$(PREFIX)/share/awk`.

To invoke the script run

```
$ awk -f abs2rel.awk \
    -v RefNode=<label> \
```

```

[-v RefOnly={0|1}] \
[-v InputMode=<imode>] \
[-v OutputMode=<omode>] \
infile.mov > outfile.mov

```

It requires the specification of the reference node label, which is given by the **RefNode** parameter.

The optional parameter **RefOnly** determines whether the output is transformed in a true reference frame, or it is simply re-oriented according to the orientation of the reference node.

The command also allows to specify the input and output formats of the orientations, under the assumption that rotations of all nodes are generated with the same format. The optional parameters **imode** and **omode** can be:

- **euler123** for the Euler-like angles, in degrees, according to the 123 sequence;
- **euler313** for the Euler-like angles, in degrees, according to the 313 sequence;
- **euler321** for the Euler-like angles, in degrees, according to the 321 sequence;
- **vector** for the orientation vector, in radians;
- **matrix** for the orientation matrix.

When **RefOnly** is set to 0 (the default), this script basically computes the configuration of each node as function of that of the reference one, namely

$$\mathbf{R}'_n = \mathbf{R}_0^T \mathbf{R}_n \quad (\text{C.1})$$

$$\mathbf{x}'_n = \mathbf{R}_0^T (\mathbf{x}_n - \mathbf{x}_0) \quad (\text{C.2})$$

$$\boldsymbol{\omega}'_n = \mathbf{R}_0^T (\boldsymbol{\omega}_n - \boldsymbol{\omega}_0) \quad (\text{C.3})$$

$$\dot{\mathbf{x}}'_n = \mathbf{R}_0^T (\dot{\mathbf{x}}_n - \dot{\mathbf{x}}_0 - \boldsymbol{\omega}_0 \times (\mathbf{x}_n - \mathbf{x}_0)) \quad (\text{C.4})$$

$$\dot{\boldsymbol{\omega}}'_n = \mathbf{R}_0^T (\dot{\boldsymbol{\omega}}_n - \dot{\boldsymbol{\omega}}_0 - \boldsymbol{\omega}_0 \times (\boldsymbol{\omega}_n - \boldsymbol{\omega}_0)) \quad (\text{C.5})$$

$$\ddot{\mathbf{x}}'_n = \mathbf{R}_0^T (\ddot{\mathbf{x}}_n - \ddot{\mathbf{x}}_0 - \dot{\boldsymbol{\omega}}_0 \times (\mathbf{x}_n - \mathbf{x}_0) - \boldsymbol{\omega}_0 \times \boldsymbol{\omega}_0 \times (\mathbf{x}_n - \mathbf{x}_0) - 2\boldsymbol{\omega}_0 \times (\dot{\mathbf{x}}_n - \dot{\mathbf{x}}_0)) \quad (\text{C.6})$$

where subscript 0 refers to the reference node, subscript n refers to the generic n -th node, and the prime refers to the relative values. Accelerations are only computed when present in the `.mov` file for both nodes.

When **RefOnly** is set to 1, the script simply re-orientates the velocity and the acceleration of each node with respect to that of the reference node, namely

$$\mathbf{R}'_n = \mathbf{R}_0^T \mathbf{R}_n \quad (\text{C.7})$$

$$\mathbf{x}'_n = \mathbf{R}_0^T (\mathbf{x}_n - \mathbf{x}_0) \quad (\text{C.8})$$

$$\boldsymbol{\omega}'_n = \mathbf{R}_0^T \boldsymbol{\omega}_n \quad (\text{C.9})$$

$$\dot{\mathbf{x}}'_n = \mathbf{R}_0^T \dot{\mathbf{x}}_n \quad (\text{C.10})$$

$$\dot{\boldsymbol{\omega}}'_n = \mathbf{R}_0^T \dot{\boldsymbol{\omega}}_n \quad (\text{C.11})$$

$$\ddot{\mathbf{x}}'_n = \mathbf{R}_0^T \ddot{\mathbf{x}}_n \quad (\text{C.12})$$

Appendix D

Log File Format

MBDyn by default writes a file with `.log` extension that contains generic, one time info about the analysis and the model.

This file is experimental, since it is mainly intended to be used by external software, like post-processing tools, to extract or somehow infer information about the model, without the complexity of parsing the entire input.

It is used, for example, by the script `mbdyn2easyanim.sh`, along with the corresponding `.mov` output file, to generate model files for `EasyAnimm`, as illustrated in Section C.1.

D.1 Generic Format

The format is typically made of a keyword, starting in the first column, made of alphanumeric chars (may include spaces, though), terminated by a colon, and possibly followed by arbitrary, context-dependent data. Data may span multiple rows, where a continuation row is marked by a blank in the first column.

D.2 Model Description Entries

Some of the keywords mark model entities, described in the following.

D.2.1 `aero0`

The `aerodynamic body`.

```
aero0: <label> <node_label>
      (Vec3) <trail_left> (Vec3) <lead_left>
      (Vec3) <trail_right> (Vec3) <lead_right>
```

The points are expressed as 3D vectors (`Vec3`), whose origin and orientation is expressed in the reference frame of the node. All data is in on one line, without continuation.

D.2.2 `aero2`

The two-node `aerodynamic beam`.

```
aero2: <label>
      <node_1_label>
```

```

(Vec3) <trail_left> (Vec3) <lead_left>
<node_2_label>
(Vec3) <trail_right> (Vec3) <lead_right>

```

The points are expressed as 3D vectors (**Vec3**), whose origin and orientation is expressed in the reference frame of the respective nodes. All data is in on one line, without continuation.

D.2.3 **aero3**

The three-node **aerodynamic beam**.

```

aero3: <label>
<node_1_label>
(Vec3) <trail_left> (Vec3) <lead_left>
<node_2_label>
(Vec3) <trail_center> (Vec3) <lead_center>
<node_3_label>
(Vec3) <trail_right> (Vec3) <lead_right>

```

The points are expressed as 3D vectors (**Vec3**), whose origin and orientation is expressed in the reference frame of the respective nodes. All data is in on one line, without continuation.

D.2.4 **axial rotation**

The **axialrotation** joint

```

axialrotation: <label>
<node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
<node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>

```

The **position_*** vectors indicates the location of the joint in the reference frame of the respective node, while the **orientation_*** matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.5 **beam2**

The two-node **beam**.

```

beam2: <label>
<node_1_label> (Vec3) <offset_1>
<node_2_label> (Vec3) <offset_2>

```

The beam **label** is followed by the label and the offset of each node. All data is in on one line, without continuation.

D.2.6 **beam3**

The three-node **beam**.

```

beam3: <label>
<node_1_label> (Vec3) <offset_1>
<node_2_label> (Vec3) <offset_2>
<node_3_label> (Vec3) <offset_3>

```

The beam **label** is followed by the label and the offset of each node. All data is in on one line, without continuation.

D.2.7 brake

The **brake** joint

```
brake: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The **position_*** vectors indicates the location of the joint in the reference frame of the respective node, while the **orientation_*** matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.8 clamp

The **clamp** joint

```
clamp: <label>
  <node_label> (Vec3) <position> (Mat3x3) <orientation>
  <node_label> (Vec3) <position> (Mat3x3) <orientation>
```

The format is quite obscure; the **position** and the **orientation** are repeated twice. Moreover, the **position** is always a vector of zeros, and the **orientation** is the identity matrix. Basically, the location of the clamp is assumed to be that of the node. All data is in on one line, without continuation.

D.2.9 deformable joints

The **deformable hinge**, the **deformable displacement joint** and the **deformable joint**, including the **invariant** versions, where defined

```
<joint_type>: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>

<joint_type> ::=
  [ invariant ] deformable { hinge | [ displacement ] joint }
```

The **position_*** vectors indicates the location of the joint in the reference frame of the respective node, while the **orientation_*** matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.10 distance

The **distance** joint

```
distance: <label>
  <node_1_label> (Vec3) <offset_1>
  <node_2_label> (Vec3) <offset_2>
```

The **label** of the joint, followed by the label of each node and the offset of the respective joint extremity, in the reference frame of the node. All data is in on one line, without continuation. Both **distance** and **distance with offset** joints are logged in this form.

D.2.11 gimbal rotation

The `gimbal rotation` joint

```
gimbalrotation: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors indicates the location of the joint in the reference frame of the respective node, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.12 inline

The `inline` joint

```
inline: <label>
  <node_1_label> (Vec3)<position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3)<position_2> (Mat3x3) <orientation_2>
```

The `label` of the joint, followed by the label of the node that carries the reference line, the reference point `position_1` on the line and the orientation `orientation_1` of the line, such that axis 3 is aligned with the line. The second node label and the position of the point on the plane follow; `orientation_2` is the identity matrix. All data is in on one line, without continuation.

D.2.13 inplane

The `inplane` joint

```
inplane: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `label` of the joint, followed by the label of the node that carries the reference plane, the reference point `position_1` on the plane and the orientation `orientation_1` of the plane, such that axis 3 is normal to the plane. The second node label and the position of the point on the plane follow; `orientation_2` is the identity matrix. All data is in on one line, without continuation.

D.2.14 prismatic

The `prismatic` joint

```
prismatic: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors are zero, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective nodes. All data is in on one line, without continuation.

D.2.15 relative frame structural node

The `relative frame` variant of `dummy` structural nodes is logged separately, as it cannot be intermixed with regular structural nodes and the `offset` variant of `dummy` structural nodes. The syntax is the same of the `structural node` (Section D.2.22), namely

```
relative frame structural node: <label> (Vec3) <X> <orientation_description>
```

D.2.16 revolute hinge

The `revolute hinge` joint

```
revolutehinge: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors indicates the location of the joint in the reference frame of the respective node, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.17 revolute rotation

The `revolute rotation` joint

```
revoluterotation: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors indicates the location of the joint in the reference frame of the respective node, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.18 rod

The `rod` joint

```
rod: <label>
  <node_1_label> (Vec3) <offset_1>
  <node_2_label> (Vec3) <offset_2>
```

The `label` of the joint, followed by the label of each node and the offset of the respective joint extremity, in the reference frame of the node. All data is in on one line, without continuation. Both `rod` and `rod with offset` joints are logged in this form.

D.2.19 spherical hinge

The `spherical hinge` joint

```
sphericalhinge: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors indicates the location of the joint in the reference frame of the respective node, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.20 spherical pin

The `spherical pin` joint

```
sphericalpin: <label>
  <node_label> (Vec3) <position> (Mat3x3) <orientation>
  <node_label> (Vec3) <position> (Mat3x3) <orientation>
```

The format is quite obscure; the `position` and the `orientation` are repeated twice. The vector `position` indicates the location of the joint, while the matrix `orientation` is the identity matrix. All data is in on one line, without continuation.

D.2.21 structural forces

The structural forces elements

```
<force_type>: <force_label>
  <node1_label> (Vec3)<arm1>
  [ <node2_label> (Vec3)<arm2> ]

<force_type> ::= structural [ internal ] { absolute | follower } { force | couple }
```

When `internal` is present `node2_label` and `arm2` must be present as well, otherwise they must be omitted.

D.2.22 structural node

The `structural node`

```
structural node: <label> (Vec3) <X> <orientation_description>

<orientation_description> ::=
  { euler123 (Vec3) <euler123_angles>
    | euler313 (Vec3) <euler313_angles>
    | euler321 (Vec3) <euler321_angles>
    | phi (Vec3) <orientation_vector>
    | mat (Mat3x3) <orientation_matrix> }
```

The `label` of the node, the position `X` and the `orientation` parameters that express the orientation are given. When `orientation` is any of `euler*` parametrizations, the angles are in degrees.

D.2.23 universal hinge

The `universal hinge` joint

```
universalhinge: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors indicates the location of the joint in the reference frame of the respective node, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.2.24 universal pin

The `universal pin` joint

```
universalpin: <label>
  <node_label> (Vec3) <position> (Mat3x3) <orientation>
  <node_label> (Vec3) <relative_position> (Mat3x3) <relative_orientation>
```

The `position` vector and the `orientation` matrix indicate the location and the orientation of the joint in the global reference frame, while the `relative_position` vector and the `relative_orientation` matrix indicate the location and the orientation of the joint in the reference frame of the node. All data is in on one line, without continuation.

D.2.25 universal rotation

The `universal rotation` joint

```
universalrotation: <label>
  <node_1_label> (Vec3) <position_1> (Mat3x3) <orientation_1>
  <node_2_label> (Vec3) <position_2> (Mat3x3) <orientation_2>
```

The `position_*` vectors are zero, while the `orientation_*` matrices indicate the orientation of the joint in the reference frame of the respective node. All data is in on one line, without continuation.

D.3 Analysis Description Entries

Some of the keywords mark analysis entities, described in the following.

D.3.1 inertia

An aggregate inertia.

```
inertia: <label> (<name>)
  mass:          (Scalar) <mass>
  J:             (Mat3x3) <J>
  Xcg:           (Vec3) <x_cg>
  Jcg:           (Mat3x3) <J_cg>
  Vcg:           (Vec3) <v_cg>
  Wcg:           (Vec3) <omega_cg>
  Xcg-X:        (Vec3) <relative_x_cg>
  R^T*(Xcg-X): (Vec3) <relative_orientation_x_cg>
  Rp:            (Vec3) <principal_axes_orientation_matrix>
  Thetap:       (Vec3) <principal_axes_orientation_vector>
  Jp:            (Mat3x3) <principal_inertia_matrix_cg>
```

D.3.2 output frequency

An indication of how often output occurs with respect to time steps.

```
output frequency: <how_often>
```

where `how_often` can be:

- `custom` if the output is specially crafted by the `output meter` statement (Section 5.2.7);
- a positive integer if the output occurs every `how_often` time steps, as dictated by the `output frequency` statement (Section 5.2.6).

D.3.3 struct node dofs

The list of structural nodes (excluding dummy nodes).

```
struct node dofs: <idx> [...]
```

where `idx` is the off-by-one index of the location of the first state associated to a node (for example, $X(\text{idx} + i)$, with $i = 1, 2, 3$, is the component of the position along axis i).

D.3.4 struct node momentum dofs

The list of structural nodes (excluding static and dummy nodes).

```
struct node momentum dofs: <idx> [...]
```

where `idx` is the off-by-one index of the location of the first state associated to a node's momentum and momenta moment (for example, $X(\text{idx} + i)$, with $i = 1, 2, 3$, is the component of the position along axis i).

D.3.5 struct node labels

The list of structural nodes labels (excluding dummy nodes).

```
struct node label: <label> [...]
```

Appendix E

Frequently Asked Questions

This section contains answers to frequently asked questions related to input and output formatting, and execution control. It is an attempt to fill the knowledge gap between the user's needs and the description of each input statement. In fact, this manual assumes that the user already knows what statement is required to solve a given problem or to model a given system, so that all it is missing is the details about the syntax. This is typically not true, especially for inexperienced users, whose main problem consists in understanding how to solve a problem, or model a system. Of course an input manual is not the right document to start with but, better than nothing, this section is trying to fill this gap, allowing to approach the manual "the other way round".

If one's question is listed below, then it's likely that the answer will point to reading the right sections of the manual.

E.1 Input

E.1.1 What is the exact syntax of element X ?

The exact syntax of each input card is illustrated in the input manual.

The input manual is regularly updated, but omissions may occur, and outdated stuff and bugs may always slip in. Please feel free to notify errors and submit patches, if you think there is anything wrong in it, using the `mbdyn-users@mbdyn.org` mailing list (you need to subscribe first; instructions can be obtained on the website <http://www.aero.polimi.it/mbdyn/>).

E.1.2 What element should I use to model Y ?

To answer this question, the complement of the input manual, namely a modeling manual, is required. Unfortunately, such document does not exist, and it is not even foreseen. Right now, modeling style and capabilities are grossly addressed in the tutorials; for specific needs you should ask on the `mbdyn-users@mbdyn.org` mailing list (you need to subscribe first; instructions can be obtained on the web site <http://www.aero.polimi.it/mbdyn/>).

E.2 Output

E.2.1 How can I reduce the amount of output?

There are different possibilities: one can only output a portion of the results, or produce output only at selected time steps, or both.

To selectively produce output there are different means:

- **by category:** use the `default output` statement (Section 5.2.5) to indicate what entities should be output. By default, all entities output data, so the suggested use is:

```
default output: none, structural nodes, beams;
```

so that the first value, `none`, turns all output off, while the following values re-enable the desired ones only.

- **entity by entity, directly:** all entities allow an output specification as the last argument (nodes: Section 6; elements: Section 8); so, at the end of an entity's card just put

```
..., output, no;    # to disable output
..., output, yes;   # to enable output
```

- **entity by entity, deferred:** all entities allow to re-enable output after they have been instantiated by using the `output` statement (nodes: Section 6.6.1; elements: Section 8.17.4).

So a typical strategy is:

- use `default output` to select those types that must be on or off; also disable the types of those entities that must be partially on and partially off;
- use the direct option to respectively set the output of those entities that must be absolutely on or off; this overrides the `default output` value.
- use the `output` statement to enable the output of those entities that must be on for that simulation only; this overrides both the `default output` and the direct values.

The direct option goes in the model, so it may be hidden in an include file of a component that is used in different models; its value should not be modified directly. On the contrary, the `output` statement can appear anywhere inside the node/element blocks, so the same subcomponents can have different output behavior by using different `output` statements.

Example.

```
# ...
begin: control data;
  default output: none, joints;
  structural nodes: 2;
  joints: 2;
end: control data;
begin: nodes;
  structural: 1, static,
    null, eye, null, null,
  output, no; # this is the ground; never output it
```

```

    structural: 2, dynamic,
      null, eye, null, null;
      # defer decision about output

    # uncomment to output
    # output: structural, 2;
end: nodes;
begin: elements;
  joint: 1, clamp, 1, node, node,
    output, no; # this is the ground; never output it
  joint: 2, revolute hinge,
    1, null,
      hinge, eye,
    2, null,
      hinge, eye;
    # this is output by "default output"

    # make sure that no matter what "default output"
    # is set, this joint is output
    output: joint, 2;
end: elements;

```

To reduce output in time, one can use the `output meter` statement (Section 5.2.7). It consists in a drive that causes output to be produced only when its value is not 0.

For example, to produce output only after a certain time, use

```
output meter: string, "Time > 10.";
```

to produce output every 5 time steps starting from 0., use

```
output meter: meter, 0., forever, steps, 5;
```

to produce output only when a certain variable reaches a certain value, use

```
output meter: string, "Omega > 99.";
```

if Ω is the third component of the angular velocity of a certain node (say, node 1000), later on, after the node is defined, the related definition must appear:

```
set: "[dof,Omega,1000,structural,6,differential]
```

i.e. the sixth component of the derivative state of node 1000, the angular velocity about axis 3, is assigned to the variable Ω . Any time Ω is evaluated in a string expression, its value gets updated.

E.3 Execution Debugging

E.3.1 How can I find out why the iterative solution of a nonlinear problem does not converge?

One needs to:

1. *find out the equations whose residual does not converge to zero.*

The residual is printed using the statement

```
output: residual;
```

in the problem control block described in Section 4.1.4 of Chapter 4. Each coefficient is preceded by its number and followed by a brief description of which node/element instantiated the equation and what its purpose is, if available; for example, the output

```
Eq    1:      0 ModalStructNode(1): linear velocity definition vx
Eq    2:      0 ModalStructNode(1): linear velocity definition vy
Eq    3:      0 ModalStructNode(1): linear velocity definition vz
Eq    4:      0 ModalStructNode(1): angular velocity definition wx
Eq    5:      0 ModalStructNode(1): angular velocity definition wy
Eq    6:      0 ModalStructNode(1): angular velocity definition wz
Eq    7:      0 ModalStructNode(1): force equilibrium Fx
Eq    8:      0 ModalStructNode(1): force equilibrium Fy
Eq    9:      0 ModalStructNode(1): force equilibrium Fz
Eq   10:      0 ModalStructNode(1): moment equilibrium Mx
Eq   11:      0 ModalStructNode(1): moment equilibrium My
Eq   12:      0 ModalStructNode(1): moment equilibrium Mz
```

corresponds to a residual whose first 12 equations refer to a modal node labeled “1” (and the residual is exactly zero);

2. *find out who instantiated the offending equation or equations.*

If this is not already indicated in the previously mentioned description, one should look up the offending equation index in the output originated by adding the statement

```
print: equation description;
```

in the **control data** block, as described in Section 5.2.2. In the above example, it would generate

```
Structural(1): 12 1->12
    1->3: linear velocity definition [vx,vy,vz]
    4->6: angular velocity definition [wx,wy,wz]
    7->9: force equilibrium [Fx,Fy,Fz]
    10->12: moment equilibrium [Mx,My,Mz]
```

3. *find out who could contribute to that equation or equations.*

If the equation was instantiated by a node, one should look at the elements connected to that node. This information is obtained by adding the statement

```
print: node connection;
```

in the **control data** block, as described in Section 5.2.2. In the above example, it would generate

```
Structural(1) connected to
    Joint(1)
    Joint(2)
```

so one should try to find out which of the connected elements is generating the offending contribution to that equation.

If the equation was instantiated by an element, usually the element itself is the sole contributor to that equation.

Bibliography

- [1] David J. Laino and A. Craig Hansen. User's guide to the wind turbine aerodynamics computer software AeroDyn. Technical report, NREL, 2002. Version 12.50.
- [2] Pierre Dupont, Vincent Hayward, Brian Armstrong, and Friedhelm Altpeter. Single state elastoplastic friction models. *IEEE Transactions on Automatic Control*, June 2002.
- [3] John C. Houbolt and George W. Brooks. Differential equations of motion for combined flapwise bending, chordwise bending, and torsion of twisted nonuniform rotor blades. TN 3905, NACA, 1957. Updated in 1975.
- [4] Marco Borri and Teodoro Merlini. A large displacement formulation for anisotropic beam analysis. *Meccanica*, 21:30–37, 1986.
- [5] Pierangelo Masarati, Massimiliano Lanz, and Paolo Mantegazza. Multistep integration of ordinary, stiff and differential-algebraic problems for multibody dynamics applications. In *XVI Congresso Nazionale AIDAA*, pages 71.1–10, Palermo, 24–28 Settembre 2001.
- [6] C.T. Kelley. *Iterative methods for linear and nonlinear equations*. SIAM, Philadelphia, PA, 1995.
- [7] P. Masarati. Direct eigenanalysis of constrained system dynamics. *Proc. IMech. E Part K: J. Multi-body Dynamics*, 223(4):335–342, 2009. doi:10.1243/14644193JMBD211.
- [8] Marco Morandini and Paolo Mantegazza. Using dense storage to solve small sparse linear systems. *ACM Trans. Math. Softw.*, 33(1, Article 5), March 2007. doi:10.1145/1206040.1206045.
- [9] I.C. Cheeseman and N.E. Bennett. The effect of ground on a helicopter rotor in forward flight. TR 3021, NASA, 1955.
- [10] Dale M. Pitt and David A. Peters. Theoretical prediction of dynamic-inflow derivatives. *Vertica*, 5:21–34, 1981.
- [11] Gian Luca Ghiringhelli, Pierangelo Masarati, and Paolo Mantegazza. A multi-body implementation of finite volume beams. *AIAA Journal*, 38(1):131–138, January 2000.
- [12] Dewey H. Hodges. Review of composite rotor blade modeling. *AIAA Journal*, 28(3):561–565, March 1990.
- [13] Vittorio Giavotto, Marco Borri, Paolo Mantegazza, Gian Luca Ghiringhelli, V. Caramaschi, G. C. Maffioli, and F. Mussi. Anisotropic beam theory and applications. *Computers & Structures*, 16(1–4):403–413, 1983.
- [14] Pierangelo Masarati. On the semi-analytical solution of beams subjected to distributed loads. In *XVI Congresso Nazionale AIDAA*, pages 72.1–10, Palermo, 24–28 Settembre 2001.

- [15] David A. Peters, Mnaouar ChouChane, and Mark Fulton. Helicopter trim with flap-lag-torsion and stall by an optimized controller. *J. Guidance*, 13(5):824–834, September-October 1990.
- [16] Stefania Gualdi, Marco Morandini, and Pierangelo Masarati. A deformable slider joint for multi-body applications. In *XVII Congresso Nazionale AIDAA*, Rome, Italy, September 15–19 2003. <http://mbdyn.aero.polimi.it/~masarati/Publications/SliderAIDA A2003.pdf>.
- [17] Pierangelo Masarati and Marco Morandini. An ideal homokinetic joint formulation for general-purpose multibody real-time simulation. *Multibody System Dynamics*, 20(3):251–270, October 2008. doi:10.1007/s11044-008-9112-8.
- [18] *The NetCDF Users Guide*, June 2006. NetCDF Version 3.6.1.