


LuaMetaTeX

Reference

Manual



December 2019
Version 2.03.6

LuaMetaT_EX

Reference

Manual

copyright : LuaT_EX development team
 : CONT_EXT development team
more info : www.luatex.org
 : contextgarden.net
version : December 31, 2019

Contents

Introduction	9
1 The internals	11
2 Differences with L^AT_EX	15
3 The original engines	19
3.1 The merged engines	19
3.1.1 The rationale	19
3.1.2 Changes from T _E X 3.1415926	19
3.1.3 Changes from ϵ -T _E X 2.2	20
3.1.4 Changes from PDFT _E X 1.40	21
3.1.5 Changes from ALEPH RC4	21
3.1.6 Changes from standard WEB2C	22
3.2 Implementation notes	22
3.2.1 Memory allocation	22
3.2.2 Sparse arrays	22
3.2.3 Simple single-character csnames	23
3.2.4 Binary file reading	23
3.2.5 Tabs and spaces	23
3.2.6 Logging	23
4 Using L^AMETAT_EX	25
4.1 Initialization	25
4.1.1 L ^A METAT _E X as a LUA interpreter	25
4.1.2 Other commandline processing	25
4.2 LUA behaviour	26
4.2.1 The LUA version	26
4.2.2 Locales	27
4.3 LUA modules	27
4.4 Testing	27
5 Basic T_EX enhancements	29
5.1 Introduction	29
5.1.1 Primitive behaviour	29
5.1.2 Experiments	29
5.1.3 Version information	30
5.2 UNICODE text support	30
5.2.1 Extended ranges	30
5.2.2 \Uchar	31
5.2.3 Extended tables	31



5.3	Attributes	32
5.3.1	Nodes	32
5.3.2	Attribute registers	32
5.3.3	Box attributes	33
5.4	LUA related primitives	34
5.4.1	\directlua	34
5.4.2	\luaescapestring	35
5.4.3	\luafunction, \luafunctioncall and \luadef	35
5.4.4	\luabytecode and \luabytecodecall	36
5.5	Catcode tables	37
5.5.1	Catcodes	37
5.5.2	\catcodetable	37
5.5.3	\initcatcodetable	37
5.5.4	\savecatcodetable	38
5.6	Tokens, commands and strings	38
5.6.1	\scantextokens	38
5.6.2	\toksapp, \tokspre, \etoksapp, \etokspre, \gtoksapp, \gtokspre, \xtoksapp, \xtokspre	38
5.6.3	\csstring, \begincsname and \lastnamedcs	38
5.6.4	\clearmarks	39
5.6.5	\alignmark and \aligntab	39
5.6.6	\letcharcode	39
5.6.7	\glet	39
5.6.8	\expanded, \immediateassignment and \immediateassigned	39
5.6.9	\ignorepars	41
5.6.10	\futureexpand, \futureexpandis, \futureexpandisap	41
5.6.11	\aftergrouped	41
5.7	Conditions	42
5.7.1	\ifabsnum and \ifabsdim	42
5.7.2	\ifcmpnum, \ifcmpdim, \ifnumval, \ifdimval, \ifchknum and \ifchkdim	42
5.7.3	\iftok and \ifcstok	43
5.7.4	\ifcondition	44
5.7.5	\orelse	45
5.7.6	\ifprotected, \frozen, \iffrozen and \ifusercmd	46
5.8	Boxes, rules and leaders	46
5.8.1	\outputbox	46
5.8.2	\vpack, \hpack and \tpack	46
5.8.3	\vsplit	46
5.8.4	Images and reused box objects	46
5.8.5	\hpack, \vpack and \tpack	48
5.8.6	\nohrule and \novrule	48
5.8.7	\gleaders	48



5.9	Languages	48
5.9.1	<code>\hyphenationmin</code>	48
5.9.2	<code>\boundary</code> , <code>\noboundary</code> , <code>\protrusionboundary</code> and <code>\wordboundary</code>	48
5.10	Control and debugging	49
5.10.1	Tracing	49
5.10.2	<code>\lastnodetype</code> , <code>\lastnodesubtype</code> , <code>\currentifttype</code> and <code>\internalcodesmode</code> .	49
5.11	Files	49
5.11.1	File syntax	49
5.11.2	Writing to file	49
5.12	Math	50
5.13	Fonts	50
5.14	Directions	50
5.14.1	Two directions	50
5.14.2	How it works	50
5.14.3	Controlling glue with <code>\breakafterdirmode</code>	52
5.14.4	Controlling parshapes with <code>\shapemode</code>	52
5.14.5	Orientations	53
5.15	Expressions	54
5.16	Nodes	54
6	Fonts	55
6.1	Introduction	55
6.2	Defining fonts	55
6.3	Virtual fonts	58
6.4	Additional \TeX commands	60
6.4.1	Font syntax	60
6.4.2	<code>\fontid</code> and <code>\setfontid</code>	61
6.4.3	<code>\noligs</code> and <code>\nokerns</code>	61
6.4.4	<code>\nospaces</code>	61
6.4.5	<code>\protrusionboundary</code>	62
6.5	The LUA font library	62
6.5.1	Introduction	62
6.5.2	Defining a font with <code>define</code> , <code>addcharacters</code> and <code>setfont</code>	63
6.5.3	Font ids: <code>id</code> , <code>max</code> and <code>current</code>	63
7	Math	65
7.1	Traditional alongside OPENTYPE	65
7.2	Unicode math characters	65
7.3	Math styles	66
7.3.1	<code>\mathstyle</code>	66
7.3.2	<code>\Ustack</code>	68
7.3.3	The new <code>\cramped ...style</code> commands	68
7.4	Math parameter settings	69
7.4.1	Many new <code>\Umath*</code> primitives	69
7.4.2	Font-based math parameters	71



7.5	Math spacing	75
7.5.1	Setting inline surrounding space with <code>\mathsurround[skip]</code>	75
7.5.2	Pairwise spacing and <code>\Umath...spacing</code> commands	76
7.5.3	Local settings	77
7.5.4	Skips around display math and <code>\mathdisplayskipmode</code>	77
7.5.5	Nolimit correction with <code>\mathnolimitsmode</code>	78
7.5.6	Controlling math italic mess with <code>\mathitalicsmode</code>	78
7.5.7	Influencing script kerning with <code>\mathscriptboxmode</code>	78
7.5.8	Forcing fixed scripts with <code>\mathscriptsmode</code>	79
7.5.9	Penalties: <code>\mathpenaltiesmode</code>	80
7.5.10	Equation spacing: <code>\matheqnogapstep</code>	80
7.6	Math constructs	81
7.6.1	Unscaled fences and <code>\mathdelimitersmode</code>	81
7.6.2	Accent handling with <code>\Umathaccent</code>	81
7.6.3	Building radicals with <code>\Uradical</code> and <code>\Uroot</code>	82
7.6.4	Super- and subscripts	82
7.6.5	Scripts on extensibles: <code>\Uunderdelimiter</code> , <code>\Uoverdelimiter</code> , <code>\Udelimiterover</code> , <code>\Udelimiterunder</code> and <code>\Uextensible</code>	83
7.6.6	Fractions and the new <code>\Uskewed</code> and <code>\Uskewedwithdelims</code>	84
7.6.7	Delimiters: <code>\Uleft</code> , <code>\Umiddle</code> and <code>\Uright</code>	84
7.6.8	Accents: <code>\mathlimitsmode</code>	85
7.7	Extracting values	85
7.7.1	Codes and using <code>\Umathcode</code> , <code>\Umathcharclass</code> , <code>\Umathcharfam</code> and <code>\Umathcharslot</code>	85
7.7.2	Last lines and <code>\predisplaygapfactor</code>	86
7.8	Math mode	86
7.8.1	Verbose versions of single-character math commands like <code>\Usuperscript</code> and <code>\Usubscript</code>	86
7.8.2	Script commands <code>\Unosuperscript</code> and <code>\Unosubscript</code>	86
7.8.3	Allowed math commands in non-math modes	87
7.9	Goodies	87
7.9.1	Flattening: <code>\mathflattenmode</code>	87
7.9.2	Less Tracing	88
8	Nodes	89
8.1	LUA node representation	89
8.2	Main text nodes	90
8.2.1	<code>hlist</code> and <code>vlist</code> nodes	91
8.2.2	<code>rule</code> nodes	91
8.2.3	<code>ins</code> nodes	92
8.2.4	<code>mark</code> nodes	92
8.2.5	<code>adjust</code> nodes	93
8.2.6	<code>disc</code> nodes	93
8.2.7	<code>math</code> nodes	94
8.2.8	<code>glue</code> nodes	94
8.2.9	<code>glue_spec</code> nodes	95
8.2.10	<code>kern</code> nodes	95



8.2.11	penalty nodes	95
8.2.12	glyph nodes	96
8.2.13	boundary nodes	97
8.2.14	local_par nodes	97
8.2.15	dir nodes	98
8.2.16	marginkern nodes	98
8.2.17	Whatsits	98
8.2.18	Math noads	98
8.3	The node library	103
8.3.1	Introduction	103
8.3.2	Housekeeping	104
8.3.3	Manipulating lists	106
8.3.4	Glue handling	110
8.3.5	Attribute handling	110
8.3.6	Glyph handling	112
8.3.7	Packaging	114
8.3.8	Math	116
8.4	Two access models	117
8.5	Normalization	122
8.6	Properties	122
9	LUA callbacks	127
9.1	Registering callbacks	127
9.2	File related callbacks	127
9.2.1	find_read_file	128
9.2.2	find_data_file	128
9.2.3	find_format_file	128
9.2.4	open_read_file	128
9.3	Data processing callbacks	129
9.3.1	process_jobname	129
9.4	Node list processing callbacks	130
9.4.1	contribute_filter	130
9.4.2	buildpage_filter	130
9.4.3	build_page_insert	131
9.4.4	pre_linebreak_filter	131
9.4.5	linebreak_filter	132
9.4.6	append_to_vlist_filter	132
9.4.7	post_linebreak_filter	133
9.4.8	hpack_filter	133
9.4.9	vpack_filter	133
9.4.10	hpack_quality	133
9.4.11	vpack_quality	134
9.4.12	process_rule	134
9.4.13	pre_output_filter	134
9.4.14	hyphenate	134
9.4.15	ligaturing	135
9.4.16	kerning	135



9.4.17	insert_local_par	135
9.4.18	mlist_to_hlist	135
9.5	Information reporting callbacks	136
9.5.1	pre_dump	136
9.5.2	start_run	136
9.5.3	stop_run	136
9.5.4	show_error_hook	136
9.5.5	show_error_message	137
9.5.6	show_lua_error_hook	137
9.5.7	start_file	137
9.5.8	stop_file	137
9.5.9	wrapup_run	137
9.6	Font-related callbacks	137
9.6.1	define_font	137
10	The TeX related libraries	139
10.1	The lua library	139
10.1.1	Version information	139
10.1.2	Table allocators	139
10.1.3	Bytecode registers	139
10.1.4	Chunk name registers	140
10.1.5	Introspection	140
10.2	The status library	140
10.3	The tex library	142
10.3.1	Introduction	142
10.3.2	Internal parameter values, set and get	142
10.3.3	Convert commands	146
10.3.4	Last item commands	146
10.3.5	Accessing registers: set*, get* and is*	146
10.3.6	Character code registers: [get set]*code[s]	148
10.3.7	Box registers: [get set]box	149
10.3.8	triggerbuildpage	150
10.3.9	splitbox	150
10.3.10	Accessing math parameters: [get set]math	150
10.3.11	Special list heads: [get set]list	151
10.3.12	Semantic nest levels: getnest and ptr	152
10.3.13	Print functions	153
10.3.14	Helper functions	155
10.3.15	Functions for dealing with primitives	158
10.3.16	Core functionality interfaces	162
10.3.17	Randomizers	164
10.3.18	Functions related to synctex	164
10.4	The texconfig table	165



10.5	The texio library	166
10.5.1	write	166
10.5.2	write_nl	166
10.5.3	setescape	166
10.5.4	closeinput	166
10.6	The token library	166
10.6.1	The scanner	166
10.6.2	Picking up one token	169
10.6.3	Creating tokens	169
10.6.4	Macros	170
10.6.5	Pushing back	171
10.6.6	Nota bene	171
11	The METAPOST library mplib	173
11.1	Process management	173
11.1.1	new	173
11.1.2	statistics	174
11.1.3	execute	174
11.1.4	finish	175
11.2	The end result	175
11.2.1	fill	176
11.2.2	outline	176
11.2.3	text	176
11.2.4	special	177
11.2.5	start_bounds, start_clip	177
11.3	Subsidiary table formats	177
11.3.1	Paths and pens	177
11.3.2	Colors	178
11.3.3	Transforms	178
11.3.4	Dashes	178
11.3.5	Pens and pen_info	178
11.4	Acessors	179
11.4.1	Character size information	179
11.4.2	get_[boolean numeric string path]	179
12	The PDF related libraries	181
12.1	The pdfe library	181
12.1.1	Introduction	181
12.1.2	open, openfile, new, getstatus, close, unencrypt	181
12.1.3	getsize, getversion, getnofobjects, getnofpages	182
12.1.4	get[catalog trailer info]	182
12.1.5	getpage, getbox	182
12.1.6	get[string integer number boolean name]	183
12.1.7	get[from][dictionary array stream]	183
12.1.8	[open close readfrom whole]stream	183
12.1.9	getfrom[dictionary array]	184
12.1.10	[dictionary array]totable	184



12.1.11	getfromreference	184
12.2	Memory streams	185
12.3	The pdfscanner library	185
13	Extra libraries	187
13.1	Introduction	187
13.2	File and string readers: fio and type sio	187
13.3	md5	187
13.4	sha2	188
13.5	xzip	188
13.6	xmath	188
13.7	xcomplex	190
13.8	lfs	191
13.9	pngdecode	192
13.10	basexx	192
13.11	Multibyte string functions	192
13.12	Extra os library functions	193
13.13	The lua library functions	194
14	Primitives aka commands	195
14.1	Introduction	195
14.2	Languages	195
14.3	Fonts	195
14.4	Math	196
14.5	Boxes and directions	199
14.6	Scanning	200
14.7	Typesetting	201
14.8	LUA	202
14.9	Management	202
14.10	Miscellaneous	202
Topics		203
Primitives		207
Callbacks		213
Nodes		215
Libraries		217
Statistics		223



Introduction

Around 2005 we started the L^AT_EX projects and it took about a decade to reach a state where we could consider the experiments to have reached a stable state. Pretty soon L^AT_EX could be used in production, even if some of the interfaces evolved, but CON_TE_XT was kept in sync so that was not really a problem. In 2018 the functionality was more or less frozen. Of course we might add some features in due time but nothing fundamental will change as we consider version 1.10 to be reasonable feature complete. Among the reasons is that this engine is now used outside CON_TE_XT too which means that we cannot simply change much without affecting other macro packages.

However, in reaching that state some decisions were delayed because they didn't go well with a current stable version. This is why at the 2018 CON_TE_XT meeting those present agreed that we could move on with a follow up tagged METAT_EX, a name we already had in mind for a while, but as L^AU^A is an important component, it got expanded to LUAMETAT_EX. This follow up is a lightweight companion to L^AT_EX that will be maintained alongside. More about the reasons for this follow up as well as the philosophy behind it can be found in the document(s) describing the development. During L^AT_EX development I kept track of what happened in a series of documents, parts of which were published as articles in user group journals, but all are in the CON_TE_XT distribution. I did the same with the development of LUAMETAT_EX.

The LUAMETAT_EX engine is, as said, a lightweight version of L^AT_EX, that for now targets CON_TE_XT. We will use it for possibly drastic experiments but without affecting L^AT_EX. As we can easily adapt CON_TE_XT to support both, no other macro package will be harmed when (for instance) interfaces change as part of an experiment. Of course, when we consider something to be useful, it can be ported back to L^AT_EX, but only when there are good reasons for doing so and when no compatibility issues are involved. When considering this follow up one consideration was that a lean and mean version with an extension mechanism is a bit closer to original T_EX. Of course, because we also have new primitives, this is not entirely true. The move to L^AU^A already meant that some aspects, especially system dependent ones, no longer made sense and therefore had consequences for the interface at the system level.

This manual currently has quite a bit of overlap with the L^AT_EX manual but some chapters are removed, others added and the rest has been (and will be further) adapted. It also discusses the (main) differences. Some of the new primitives or functions that show up in LUAMETAT_EX might show up in L^AT_EX at some point, others might not, so don't take this manual as reference for L^AT_EX! For now it is an experimental engine in which we can change things at will but with CON_TE_XT in tandem so that this macro package will keep working.

For CON_TE_XT users the LUAMETAT_EX engine will become the default. The CON_TE_XT variant for this engine is tagged LMTX. The pair can be used in production, just as with L^AT_EX and MKIV. In fact, most users will probably not really notice the difference. In some cases there will be a drop in performance, due to more work being delegated to L^AU^A, but on the average performance will be better, also due to some changes below the hood of the engine.

As this follow up is closely related to CON_TE_XT development, and because we expect stock L^AT_EX to be used outside the CON_TE_XT proper, there will be no special mailing list nor coverage (or pollution) on the L^AT_EX related mailing lists. We have the CON_TE_XT mailing lists for that. In due time the source code will be part of the regular CON_TE_XT distribution.



This manual sometimes refers to L^AT_EX, especially when we talk of features common to both engine, as well as to L^AM^ET_EX, when it is more specific to the follow up. A substantial amount of time went into the transition and more will go in, so if you want to complain about L^AM^ET_EX, don't bother me. Of course, if you really need professional support with these engines (or T_EX in general), you can always consider contacting the developers.

Hans Hagen

Version : December 31, 2019

L^AM^ET_EX : luametateX 2.036 / 20191231

C^ON^TE^XT : M_kIV 2019.12.30 19:27

L^AT_EX Team : Hans Hagen, Hartmut Henkel, Taco Hoekwater, Luigi Scarso

remark: L^AM^ET_EX development is mostly done by Hans Hagen and Alan Braslau, who love playing with the three languages involved. Testing is done by C^ON^TE^XT developers and users. Many thanks for their patience!

remark: When there are non-intrusive features that also make sense in L^AT_EX, these will be applied in the experimental branch first, so that there is no interference with the stable release.

remark: Most C^ON^TE^XT users seem always willing to keep up with the latest versions which means that L^AM^ET_EX is tested well. We can therefore safely claim that end of 2019 the code has become quite stable. There are no complaints about performance (on my laptop this manual compiles at 22.5 pps with L^AM^ET_EX versus 20.7 pps for the L^AT_EX manual with M_kIV). Probably no one notices it, but memory consumption stepwise got reduced too. And ... the binary is still below 3 MegaBytes on all platforms.



1 The internals

This is a reference manual, not a tutorial. This means that we discuss changes relative to traditional \TeX and also present new functionality. As a consequence we will refer to concepts that we assume to be known or that might be explained later. Because the \LaTeX and \LaTeX engines open up \TeX there's suddenly quite some more to explain, especially about the way a (to be) typeset stream moves through the machinery. However, discussing all that in detail makes not much sense, because deep knowledge is only relevant for those who write code not possible with regular \TeX and who are already familiar with these internals (or willing to spend time on figuring it out).

So, the average user doesn't need to know much about what is in this manual. For instance fonts and languages are normally dealt with in the macro package that you use. Messing around with node lists is also often not really needed at the user level. If you do mess around, you'd better know what you're dealing with. Reading "The \TeX Book" by Donald Knuth is a good investment of time then also because it's good to know where it all started. A more summarizing overview is given by "The \TeX by Topic" by Victor Eijkhout. You might want to peek in "The ϵ - \TeX manual" too.

But ... if you're here because of \Lua , then all you need to know is that you can call it from within a run. If you want to learn the language, just read the well written \Lua book. The macro package that you use probably will provide a few wrapper mechanisms but the basic `\directlua` command that does the job is:

```
\directlua{tex.print("Hi there")}
```

You can put code between curly braces but if it's a lot you can also put it in a file and load that file with the usual \Lua commands. If you don't know what this means, you definitely need to have a look at the \Lua book first.

If you still decide to read on, then it's good to know what nodes are, so we do a quick introduction here. If you input this text:

Hi There

eventually we will get a linked lists of nodes, which in ASCII art looks like:

```
H <=> i <=> [glue] <=> T <=> h <=> e <=> r <=> e
```

When we have a paragraph, we actually get something:

```
[localpar] <=> H <=> i <=> [glue] <=> T <=> h <=> e <=> r <=> e <=> [glue]
```

Each character becomes a so called glyph node, a record with properties like the current font, the character code and the current language. Spaces become glue nodes. There are many node types that we will discuss later. Each node points back to a previous node or next node, given that these exist. Sometimes multiple characters are represented by one glyphs, so one can also get:

```
[localpar] <=> H <=> i <=> [glue] <=> Th <=> e <=> r <=> e <=> [glue]
```



And maybe some characters get positioned relative to each other, so we might see:

```
[localpar] <=> H <=> [kern] <=> i <=> [glue] <=> Th <=> e <=> r <=> e <=> [glue]
```

It's also good to know beforehand that $\text{T}_{\text{E}}\text{X}$ is basically centered around creating paragraphs and pages. The par builder takes a list and breaks it into lines. At some point horizontal blobs are wrapped into vertical ones. Lines are so called boxes and can be separated by glue, penalties and more. The page builder accumulates lines and when feasible triggers an output routine that will take the list so far. Constructing the actual page is not part of $\text{T}_{\text{E}}\text{X}$ but done using primitives that permit manipulation of boxes. The result is handled back to $\text{T}_{\text{E}}\text{X}$ and flushed to a (often PDF) file.

The $\text{LUA}_{\text{T}_{\text{E}}\text{X}}$ engine provides hooks for LUA code at nearly every reasonable point in the process: collecting content, hyphenating, applying font features, breaking into lines, etc. This means that you can overload $\text{T}_{\text{E}}\text{X}$'s natural behaviour, which still is the benchmark. When we refer to 'callbacks' we means these hooks. The $\text{T}_{\text{E}}\text{X}$ engine itself is pretty well optimized but when you kick in much LUA code, you will notices that performance drops. Don't blame and bother the authors with performance issues. In $\text{CON}_{\text{T}_{\text{E}}\text{X}}\text{T}$ over 50% of the time can be spent in LUA, but so far we didn't get many complaints about efficiency.

Where plain $\text{T}_{\text{E}}\text{X}$ is basically a basic framework for writing a specific style, macro packages like $\text{CON}_{\text{T}_{\text{E}}\text{X}}\text{T}$ and $\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$ provide the user a whole lot of additional tools to make documents look good. They hide the dirty details of font management, language demands, turning structure into typeset results, wrapping pages, including images, and so on. You should be aware of the fact that when you hook in your own code to manipulate lists, this can interfere with the macro package that you use. Each successive step expects a certain result and if you mess around to much, the engine eventually might bark and quit. It can even crash, because testing everywhere for what users can do wrong is no real option.

When you read about nodes in the following chapters it's good to keep in mind their commands that relate to then. Here are a few:

COMMAND	NODE	EXPLANATION
<code>\hbox</code>	hlist	horizontal box
<code>\vbox</code>	vlist	vertical box with the baseline at the bottom
<code>\vtop</code>	vlist	vertical box with the baseline at the top
<code>\hskip</code>	glue	horizontal skip with optional stretch and shrink
<code>\vskip</code>	glue	vertical skip with optional stretch and shrink
<code>\kern</code>	kern	horizontal or vertical fixed skip
<code>\discretionary</code>	disc	hyphenation point (pre, post, replace)
<code>\char</code>	glyph	a character
<code>\hrule</code>	rule	a horizontal rule
<code>\vrule</code>	rule	a vertical rule
<code>\textdirection</code>	dir	a change in text direction

Text (interspersed with macros) comes from an input medium. This can be a file, token list, macro body cq. arguments, some internal quantity (like a number), LUA, etc. Macros get expanded. In the process $\text{T}_{\text{E}}\text{X}$ can enter a group. Inside the group, changes to registers get saved on a stack, and restored after leaving the group. When conditionals are encountered, another kind



of nesting happens, and again there is a stack involved. Tokens, expansion, stacks, input levels are all terms used in the next chapters. Don't worry, they lose their magic once you use T_EX a lot. You have access to most of the internals and when not, at least it is possible to query some state we're in or level we're at.

When we talk about packing it can mean two things. When T_EX has consumed some tokens that represent text the next can happen. When the text is put into a so called `\hbox` it (normally) first gets hyphenated, next ligatures are build, and finally kerns are added. Each of that stages can be overloaded using LUA code. When these three stages are finished, the dimension of the content is calculated and the box gets its width, height and depth. What happens with the box depends on what macros do with it.

The other thing that can happen is that the text starts a new paragraph. In that case some (directional) information is put in front, indentation is prepended and some skip appended at the end. Again the three stages are applied but this time, afterwards, the long line is broken into lines and the result is either added to the content of a box or to the main vertical list (the running text so to say). This is called par building. At some point T_EX decides that enough is enough and it will trigger the page builder. So, building is another concept we will encounter. Another example of a builder is the one that turns an intermediate math list into something typeset.

Wrapping something in a box is called packing. Adding something to a list is described in terms of contributing. The more complicated processes are wrapped into builders. For now this should be enough to enable you to understand the next chapters. The text is not as enlightening and entertaining as Don Knuths books, sorry.





2 Differences with L^AT_EX

As L^AMETAT_EX is a leaner and meaner L^AT_EX, this chapter will discuss what is gone. We start with the primitives that were dropped.

fonts	<code>\letterspacefont \copyfont \expandglyphsinfont \ignoreligaturesinfont \tagcode</code>
backend	<code>\dviextension \dvivariable \dvifedback \pdfextension \pdfvariable \pdffeedback \dviextension \draftmode \outputmode</code>
dimensions	<code>\pageleftoffset \pagerightoffset \pagetopoffset \pagebottomoffset \pageheight \pagewidth</code>
resources	<code>\saveboxresource \useboxresource \lastsavedboxresourceindex \saveimageresource \useimageresource \lastsavedimageresourceindex \lastsavedimageresourcepages</code>
positioning	<code>\savepos \lastxpos \lastypos</code>
directions	<code>\textdir \linedir \mathdir \pardir \pagedir \bodydir \pagedirection \bodydirection</code>
randomizer	<code>\randomseed \setrandomseed \normaldeviate \uniformdeviate</code>
utilities	<code>\synctex</code>
extensions	<code>\latelua \lateluafunction \immediate \openout \write \closeout</code>
control	<code>\suppressfontnotfounderror \suppresslongerror \suppressprimitiveerror \suppressmathparerror \suppressifcsnameerror \suppressoutererror \mathopton</code>
whatever	<code>\primitive \ifprimitive</code>
ignored	<code>\long \outer \mag</code>

The resources and positioning primitives are actually useful but can be defined as macros that (via LUA) inject nodes in the input that suit the macro package and backend. The three-letter direction primitives are gone and the numeric variants are now leading. There is no need for page and body related directions and they don't work well in L^AT_EX anyway. We only have two directions left.

The primitive related extensions were not that useful and reliable so they have been removed. There are some new variants that will be discussed later. The `\outer` and `\long` prefixes are gone as they don't make much sense nowadays and them becoming dummies opened the way to something new, again to be discussed elsewhere. I don't think that (C_{ON}T_EX_T) users will notice it. The `\suppress...` features are now default.

The `\shipout` primitive does no ship out but just erases the content of the box, if that hasn't happened already in another way.

The extension primitives relate to the backend (when not immediate) and can be implemented as part of a backend design using generic whatsits. There is only one type of whatsit now. In fact we're now closer to original T_EX with respect to the extensions.

The `img` library has been removed as it's rather bound to the backend. The `slunicode` library is also gone. There are some helpers in the string library that can be used instead and one can write additional LUA code if needed. There is no longer a pdf backend library.



In the node, tex and status library we no longer have helpers and variables that relate to the backend. The L^AT_EX engine is in principle DVI and PDF unaware. There are only generic whatsit nodes that can be used for some management related tasks. For instance you can use them to implement user nodes.

The KPSE library is no longer built-in. Because there is no backend, quite some file related callbacks could go away. The following file related callbacks remained (till now):

```
find_write_file find_data_file find_format_file
open_data_file read_data_file
```

Also callbacks related to errors stay:

```
show_error_hook show_lua_error_hook,
show_error_message show_warning_message
```

The (job) management hooks are kept:

```
process_jobname
start_run stop_run wrapup_run
pre_dump
start_file stop_file
```

Because we use a more generic whatsit model, there is a new callback:

```
show_whatsit
```

Being the core of extensibility, the typesetting callbacks of course stayed. This is what we ended up with:

```
find_log_file, find_data_file, find_format_file, open_data_file, read_data_file,
process_jobname, start_run, stop_run, define_font, pre_output_filter,
buildpage_filter, hpack_filter, vpack_filter, hyphenate, ligaturing, kerning,
pre_linebreak_filter, linebreak_filter, post_linebreak_filter,
append_to_vlist_filter, mlist_to_hlist, pre_dump, start_file, stop_file,
handle_error_hook, show_error_hook, show_lua_error_hook, show_error_message,
show_warning_message, hpack_quality, vpack_quality, insert_local_par,
contribute_filter, build_page_insert, wrapup_run, new_graf, make_extensible,
show_whatsit, terminal_input,
```

As in L^AT_EX font loading happens with the following callback. This time it really needs to be set because there is no built-in font loader.

```
define_font
```

There are all kind of subtle differences in the implementation, for instance we no longer intercept * and & as these were already replaced long ago in T_EX engines by command line options. Talking of options, only a few are left.

We took our time for reaching a stable state in L^AT_EX. Among the reasons is the fact that most was experimented with in C^ON^TE^XT. It took many man-years to decide what to keep and how



to do things. Of course there are places when things can be improved and it might happen in LUAMETAT_EX. Contrary to what is sometimes suggested, the LUAT_EX-CON_TE_XT MKIV combination (assuming matched versions) has been quite stable. It made no sense otherwise. Most CON_TE_XT functionality didn't change much at the user level. Of course there have been issues, as is natural with everything new and beta, but we have a fast update cycle.

The same is true for LUAMETAT_EX and CON_TE_XT LMTX: it can be used for production as usual and in practice CON_TE_XT users tend to use the beta releases, which proves this. Of course, if you use low level features that are experimental you're on your own. Also, as with LUAT_EX it might take many years before a long term stable is defined. The good news is that, the source code being part of the CON_TE_XT distribution, there is always a properly working, more or less long term stable, snapshot.

The error reporting subsystem has been redone a little but is still fundamentally the same. We don't really assume interactive usage but if someone uses it, it might be noticed that it is not possible to backtrack or inject something. Of course it is no big deal to implement all that in LUA if needed. It removes a system dependency and makes for a bit cleaner code.

There are new primitives too as well as some extensions to existing primitive functionality. These are described in following chapters but there might be hidden treasures in the binary. If you locate them, don't automatically assume them to stay, some might be part of experiments!





3 The original engines

3.1 The merged engines

3.1.1 The rationale

The first version of L^AT_EX, made by Hartmut after we discussed the possibility of an extension language, only had a few extra primitives and it was largely the same as PDF_T_EX. It was presented to the public in 2005. As part of the Oriental _T_EX project, Taco merged substantial parts of ALEPH into the code and some more primitives were added. Then we started more fundamental experiments. After many years, when the engine had become more stable, the decision was made to clean up the rather hybrid nature of the program. This means that some primitives were promoted to core primitives, often with a different name, and that others were removed. This also made it possible to start cleaning up the code base. In chapter 5 we discussed some new primitives, here we will cover most of the adapted ones.

During more than a decade stepwise new functionality was added and after 10 years the more of less stable version 1.0 was presented. But we continued and after some 15 years the LUAMETAT_EX follow up entered its first testing stage. But before details about the engine are discussed in successive chapters, we first summarize where we started from. Keep in mind that in LUAMETAT_EX we have a bit less than in L^AT_EX, so this section differs from the one in the L^AT_EX manual.

Besides the expected changes caused by new functionality, there are a number of not-so-expected changes. These are sometimes a side-effect of a new (conflicting) feature, or, more often than not, a change necessary to clean up the internal interfaces. These will also be mentioned.

3.1.2 Changes from _T_EX 3.1415926

Of course it all starts with traditional _T_EX. Even if we started with PDF_T_EX, most still comes from original Knuthian _T_EX. But we divert a bit.

- ▶ The current code base is written in C, not PASCAL. The original CWEB documentation is kept when possible and not wrapped in tagged comments. As a consequence instead of one large file plus change files, we now have multiple files organized in categories like tex, lua, languages, fonts, libraries, etc. There are some artifacts of the conversion to C, but these got (and get) removed stepwise. The documentation, which is actually comes from the mix of engines (via so called change files) is kept as much as possible. Of course we want to stay as close as possible to the original so that the documentation of the fundamentals behind _T_EX by Don Knuth still applies. However, because we use C, some documentation is a bit off. Also, most global variables are now collected in structures, but the original names were kept. There are lots of so called macros too.
- ▶ See chapter ?? for many small changes related to paragraph building, language handling and hyphenation. The most important change is that adding a brace group in the middle of a word (like in of{ }fice) does not prevent ligature creation. Also, the hyphenation, ligature



building and kerning has been split so that we can hook in alternative or extra code wherever we like. There are various options to control discretionary injection and related penalties are now integrated in these nodes. Language information is now bound to glyphs. The number of languages in LUAMETAT_EX is smaller than in L_UA_TE_X.

- ▶ There is no pool file, all strings are embedded during compilation. This also removed some memory constraints. We kept token and node memory management because it is convenient and efficient but parts were reimplemented in order to remove some constraints. Token memory management is largely the same.
- ▶ The specifier `plus 1 fillll` does not generate an error. The extra 'l' is simply typeset.
- ▶ The upper limit to `\endlinechar` and `\newlinechar` is 127.
- ▶ Because the backend is not built-in, the magnification (`\mag`) primitive is not doing nothing. A shipout just discards the content of the given box. The write related primitives have to be implemented in the used macro package using LUA. None of the PDF_TE_X derived primitives is present.
- ▶ There is more control over some (formerly hard-coded) math properties. In fact, there is a whole extra bit of math related code because we need to deal with OPENTYPE fonts.
- ▶ The `\outer` and `\long` prefixed are silently ignored. It is permitted to use `\par` in math.
- ▶ Because there is no font loader, a LUA variant is free to either support or not the OMEGA ofm file format. As there are hardly any such fonts it probably makes no sense.
- ▶ The lack of a backend means that some primitives related to it are not implemented. This is no big deal because it is possible to use the scanner library to implement them as needed, which depends on the macro package and backend.
- ▶ When detailed logging is enabled more detail is output with respect to what nodes are involved. This is a side effect of the core nodes having more detailed subtype information. The benefit of more detail wins from any wish to be byte compatible in the logging. One can always write additional logging in LUA.

3.1.3 Changes from ε -T_EX 2.2

Being the de-facto standard extension of course we provide the ε -T_EX features, but with a few small adaptations.

- ▶ The ε -T_EX functionality is always present and enabled so the prepended asterisk or `-etex` switch for INIT_EX is not needed.
- ▶ The T_EXXE_T extension is not present, so the primitives `\TeXxETstate`, `\beginR`, `\beginL`, `\endR` and `\endL` are missing. Instead we used the OMEGA/ALEPH approach to directionality as starting point, albeit it has been changed quite a bit, so that we're probably not that far from T_EXXE_T.
- ▶ Some of the tracing information that is output by ε -T_EX's `\tracingassigns` and `\tracingrestores` is not there. Also keep in mind that tracing doesn't involve what LUA does.
- ▶ Register management in LUAMETAT_EX uses the OMEGA/ALEPH model, so the maximum value is 65535 and the implementation uses a flat array instead of the mixed flat & sparse model from ε -T_EX.
- ▶ Because we don't use change files on top of original T_EX, the integration of ε -T_EX functionality is bit more natural, code wise.



3.1.4 Changes from PDF_{TEX} 1.40

Because we want to produce PDF the most natural starting point was the popular PDF_{TEX} program. We inherit the stable features, dropped most of the experimental code and promoted some functionality to core L_UA_TE_X functionality which in turn triggered renaming primitives. However, as the backend was dropped, not that much from PDF_{TEX} is present any more. Basically all we now inherit from PDF_{TEX} is expansion and protrusion but even that has been adapted. So don't expect L_UA_ME_TA_TE_X to be compatible.

- ▶ The experimental primitives `\ifabsnum` and `\ifabsdim` have been promoted to core primitives.
- ▶ The primitives `\ifincsname`, `\expanded` and `\quitvmode` have become core primitives.
- ▶ As the hz (expansion) and protrusion mechanism are part of the core the related primitives `\lrcode`, `\rrcode`, `\efcode`, `\leftmarginkern`, `\rightmarginkern` are promoted to core primitives. The two commands `\protrudechars` and `\adjustspacing` control these processes.
- ▶ In L_UA_ME_TA_TE_X three extra primitives can be used to overload the font specific settings: `\adjustspacingstep` (max: 100), `\adjustspacingstretch` (max: 1000) and `\adjustspacingshrink` (max: 500).
- ▶ The hz optimization code has been partially redone so that we no longer need to create extra font instances. The front- and backend have been decoupled and the glyph and kern nodes carry the used values. In L_UA_TE_X that made a more efficient generation of PDF code possible. It also resulted in much cleaner code. The backend code is gone, but of course the information is still carried around.
- ▶ When `\adjustspacing` has value 2, hz optimization will be applied to glyphs and kerns. When the value is 3, only glyphs will be treated. A value smaller than 2 disables this feature. With value of 1, font expansion is applied after T_EX's normal paragraph breaking routines have broken the paragraph into lines. In this case, line breaks are identical to standard T_EX behavior (as with PDF_{TEX}). But ... this is a left-over from the early days of PDF_{TEX} when this feature was part of a research topic. At some point level 1 might be dropped from L_UA_ME_TA_TE_X.
- ▶ When `\protrudechars` has a value larger than zero characters at the edge of a line can be made to hang out. A value of 2 will take the protrusion into account when breaking a paragraph into lines. A value of 3 will try to deal with right-to-left rendering; this is a still experimental feature.
- ▶ The pixel multiplier dimension `\pxdimen` has been inherited as core primitive.
- ▶ The primitive `\tracingfonts` is now a core primitive but doesn't relate to the backend.

3.1.5 Changes from ALEPH RC4

In L_UA_TE_X we took the 32 bit aspects and much of the directional mechanisms and merged it into the PDF_{TEX} code base as starting point for further development. Then we simplified directionality, fixed it and opened it up. In L_UA_ME_TA_TE_X not that much of the later is left. We only have two horizontal directions. Instead of vertical directions we introduce an orientation model bound to boxes.

The already reduced-to-four set of directions now only has two members: left-to-right and right-to-left. They don't do much as it is the backend that has to deal with them. When paragraphs



are constructed a change in horizontal direction is irrelevant for calculating the dimensions. So, basically most that we do is registering state and passing that on till the backend can do something with it.

Here is a summary of inherited functionality:

- ▶ The ^^ notation has been extended: after ^^^^ four hexadecimal characters are expected and after ^^^^^ six hexadecimal characters have to be given. The original T_EX interpretation is still valid for the ^^ case but the four and six variants do no backtracking, i.e. when they are not followed by the right number of hexadecimal digits they issue an error message. Because ^^^ is a normal T_EX case, we don't support the odd number of ^^^^^ either.
- ▶ Glues *immediately after* direction change commands are not legal breakpoints. There is a bit more sanity testing for the direction state.
- ▶ The placement of math formula numbers is direction aware and adapts accordingly. Boxes carry directional information but rules don't.
- ▶ There are no direction related primitives for page and body directions. The paragraph, text and math directions are specified using primitives that take a number.

3.1.6 Changes from standard WEB2C

The LUAMETAT_EX codebase is not dependent on the WEB2C framework. The interaction with the file system and TDS is up to LUA. There still might be traces but eventually the code base should be lean and mean. The METAPOST library is coded in CWEB and in order to be independent from related tools, conversion to C is done with a LUA script ran by, surprise, LUAMETAT_EX.

3.2 Implementation notes

3.2.1 Memory allocation

The single internal memory heap that traditional T_EX used for tokens and nodes is split into two separate arrays. Each of these will grow dynamically when needed. Internally a token or node is an index into these arrays. This permits for an efficient implementation and is also responsible for the performance of the core. The original documentation in T_EX The Program mostly applies!

3.2.2 Sparse arrays

The `\mathcode`, `\delcode`, `\catcode`, `\sfcode`, `\lccode` and `\uccode` (and the new `\hjcode`) tables are now sparse arrays that are implemented in C. They are no longer part of the T_EX 'equivalence table' and because each had 1.1 million entries with a few memory words each, this makes a major difference in memory usage. Performance is not really hurt by this.

The `\catcode`, `\sfcode`, `\lccode`, `\uccode` and `\hjcode` assignments don't show up when using the ε -T_EX tracing routines `\tracingassigns` and `\tracingrestores` but we don't see that as a real limitation. It also saves a lot of clutter.

A side-effect of the current implementation is that `\global` is now more expensive in terms of processing than non-global assignments but not many users will notice that.



The glyph ids within a font are also managed by means of a sparse array as glyph ids can go up to index $2^{21} - 1$ but these are never accessed directly so again users will not notice this.

3.2.3 Simple single-character csnames

Single-character commands are no longer treated specially in the internals, they are stored in the hash just like the multiletter csnames.

The code that displays control sequences explicitly checks if the length is one when it has to decide whether or not to add a trailing space.

Active characters are internally implemented as a special type of multi-letter control sequences that uses a prefix that is otherwise impossible to obtain.

3.2.4 Binary file reading

All of the internal code is changed in such a way that if one of the `read_xxx_file` callbacks is not set, then the file is read by a C function using basically the same convention as the callback: a single read into a buffer big enough to hold the entire file contents. While this uses more memory than the previous code (that mostly used `getc` calls), it can be quite a bit faster (depending on your IO subsystem). So far we never had issues with this approach.

3.2.5 Tabs and spaces

We conform to the way other \TeX engines handle trailing tabs and spaces. For decades trailing tabs and spaces (before a newline) were removed from the input but this behaviour was changed in September 2017 to only handle spaces. We are aware that this can introduce compatibility issues in existing workflows but because we don't want too many differences with upstream \TeX LIVE we just follow up on that patch (which is a functional one and not really a fix). It is up to macro packages maintainers to deal with possible compatibility issues and in \LaTeX they can do so via the callbacks that deal with reading from files.

The previous behaviour was a known side effect and (as that kind of input normally comes from generated sources) it was normally dealt with by adding a comment token to the line in case the spaces and/or tabs were intentional and to be kept. We are aware of the fact that this contradicts some of our other choices but consistency with other engines. We still stick to our view that at the log level we can (and might be) more incompatible. We already expose some more details anyway.

3.2.6 Logging

The information that goes into the log file can be different from \LaTeX , and might even differ a bit more in the future. The main reason is that inside the engine we have more granularity, which for instance means that we output subtype related information when nodes are printed. Of course we could have offered a compatibility mode but it serves no purpose. Over time there have been many subtle changes to control logs in the \TeX ecosystems so another one is bearable.

In a similar fashion, there is a bit different behaviour when \TeX expects input, which in turn is a side effect of removing the interception of `*` and `&` which made for cleaner code (quite a bit



had accumulated as side effect of continuous adaptations in the $\text{T}_{\text{E}}\text{X}$ ecosystems). There was already code that was never executed, simply as side effect of the way $\text{LUA}\text{T}_{\text{E}}\text{X}$ initializes itself (one needs to enable classes of primitives for instance).



4 Using L^AMETAT_EX

4.1 Initialization

4.1.1 L^AMETAT_EX as a LUA interpreter

Although L^AMETAT_EX is primarily meant as a T_EX engine, it can also serve as a stand alone LUA interpreter. There are two ways to make L^AMETAT_EX behave like a standalone LUA interpreter:

- if a `--luaonly` option is given on the commandline, or
- if the only non-option argument (file) on the commandline has the extension `lua` or `luc`.

In this mode, it will set LUA's `arg[0]` to the found script name, pushing preceding options in negative values and the rest of the command line in the positive values, just like the LUA interpreter does.

L^AMETAT_EX will exit immediately after executing the specified LUA script and is, in effect, a somewhat bulky stand alone LUA interpreter with a bunch of extra preloaded libraries.

When no argument is given, L^AMETAT_EX will look for a LUA file with the same name as the binary and run that one when present. This makes it possible to use the engine as a stub. For instance, in CON_TE_XT a symlink from `mtxrun` to type `luametatex` will run the `mtxrun.lua` script when present in the same path as the binary itself

4.1.2 Other commandline processing

When the L^AMETAT_EX executable starts, it looks for the `--lua` command line option. If there is no `--lua` option, the command line is interpreted in a similar fashion as the other T_EX engines. All options are accepted but only some have are understood by L^AMETAT_EX itself:

COMMANDLINE ARGUMENT	EXPLANATION
<code>--credits</code>	display credits and exit
<code>--fmt=FORMAT</code>	load the format file <code>FORMAT</code>
<code>--help</code>	display help and exit
<code>--ini</code>	be <code>iniluatex</code> , for dumping formats
<code>--jobname=STRING</code>	set the job name to <code>STRING</code>
<code>--lua=FILE</code>	load and execute a LUA initialization script
<code>--version</code>	display version and exit

There are less options than with L^AUAT_EX, because one has to deal with them in LUA anyway. There are no options to enter a safer mode or control executing programs. This can easily be achieved with a startup LUA script.

The value to use for `\jobname` is decided as follows:



- ▶ If `--jobname` is given on the command line, its argument will be the value for `\jobname`, without any changes. The argument will not be used for actual input so it need not exist. The `--jobname` switch only controls the `\jobname` setting.
- ▶ Otherwise, `\jobname` will be the name of the first file that is read from the file system, with any path components and the last extension (the part following the last `.`) stripped off.
- ▶ There is an exception to the previous point: if the command line goes into interactive mode (by starting with a command) and there are no files input via `\everyjob` either, then the `\jobname` is set to `texput` as a last resort.

Next the initialization script is loaded and executed. From within the script, the entire command line is available in the LUA table `arg`, beginning with `arg[0]`, containing the name of the executable. As consequence warnings about unrecognized options are suppressed.

Command line processing happens very early on. So early, in fact, that none of \TeX 's initializations have taken place yet. The LUA libraries that don't deal with \TeX are initialized early.

LUAMETAT \TeX allows some of the command line options to be overridden by reading values from the `texconfig` table at the end of script execution (see the description of the `texconfig` table later on in this document for more details on which ones exactly).

So let's summarize this. The handling of when is called `jobname` is a bit complex. There can be explicit names set on the command line but when not set they can be taken from the `texconfig` table.

startup filename	<code>--lua</code>	a LUA file	
startup jobname	<code>--jobname</code>	a \TeX tex	<code>texconfig.jobname</code>
startup dumpname	<code>--fmt</code>	a format file	<code>texconfig.formatname</code>

These names are initialized according to `--luaonly` or the first filename seen in the list of options. Special treatment of `&` and `*` as well as interactive startup is gone.

When we are in \TeX mode at some point the engine needs a filename, for instance for opening a log file. At that moment the `setjobname` becomes the internal one and when it has not been set which internalized to `jobname` but when not set becomes `texput`. When you see a `texput.log` file someplace on your system it normally indicates a bad run.

When running on MS WINDOWS the command line, filenames, environment variable access etc. internally uses the current code page but to the user is exposed as UTF8. Normally users won't notice this.

4.2 LUA behaviour

4.2.1 The LUA version

We currently use LUA 5.4 and will follow developments of the language but normally with some delay. Therefore the user needs to keep an eye on (subtle) differences in successive versions of the language. Here is an example of one aspect.

LUA's `tostring` function (and `string.format` may return values in scientific notation, thereby confusing the \TeX end of things when it is used as the right-hand side of an assignment to a



`\dimen` or `\count`. The output of these serializers also depend on the LUA version, so in LUA 5.3 you can get different output than from 5.2. It is best not to depend the automatic cast from string to number and vice versa as this can change in future versions.

4.2.2 Locales

In stock LUA, many things depend on the current locale. In LUAMETAT_EX, we can't do that, because it makes documents unportable. While LUAMETAT_EX is running it forces the following locale settings:

```
LC_CTYPE=C
LC_COLLATE=C
LC_NUMERIC=C
```

There is no way to change that as it would interfere badly with the often language specific conversions needed at the T_EX end.

4.3 LUA modules

Of course the regular LUA modules are present. In addition we provide the `lpeg` library by Roberto Ierusalimsky. This library is not UNICODE-aware, but interprets strings on a byte-per-byte basis. This mainly means that `lpeg.S` cannot be used with UTF8 characters encoded in more than two bytes, and thus `lpeg.S` will look for one of those two bytes when matching, not the combination of the two. The same is true for `lpeg.R`, although the latter will display an error message if used with multibyte characters. Therefore `lpeg.R('ä')` results in the message `bad argument #1 to 'R' (range must have two characters)`, since to `lpeg`, `ä` is two 'characters' (bytes), so `ä` totals three. In practice this is no real issue and with some care you can deal with UNICODE just fine.

There are some more libraries present. These are discussed on a later chapter. For instance we embed `luasocket` but contrary to LUAT_EX don't embed the related LUA code. An adapted version of `luafilesystem` is also included. There is a more extensive math library and there are libraries that deal with encryption and compression.

4.4 Testing

For development reasons you can influence the used startup date and time. By setting the `start_time` variable in the `texconfig` table; as with other variables we use the internal name there. When Universal Time is needed, set the entry `use_utc_time` in the `texconfig` table.

In CON_TE_XT we provide the command line argument `--nodates` that does a bit more than disabling dates; it avoids time dependent information in the output file for instance.





5 Basic T_EX enhancements

5.1 Introduction

5.1.1 Primitive behaviour

From day one, L^AT_EX has offered extra features compared to the superset of P^DF_T_EX, which includes ε -T_EX, and ALEPH. This has not been limited to the possibility to execute LUA code via `\directlua`, but L^AT_EX also adds functionality via new T_EX-side primitives or extensions to existing ones. The same is true for L^AMETAT_EX. Some primitives have `luatex` in their name and there will be no `luametate` variants. This is because we consider L^AMETAT_EX to be L^AT_EX²⁺.

Contrary to the L^AT_EX engine L^AMETAT_EX enables all its primitives. You can clone (a selection of) primitives with a different prefix, like:

```
\directlua { tex.enableprimitives('normal',tex.extraprimitives()) }
```

The `extraprimitives` function returns the whole list or a subset, specified by one or more keywords `core`, `tex`, `etex` or `luatex`.¹

But be aware that the curly braces may not have the proper `\catcode` assigned to them at this early time (giving a ‘Missing number’ error), so it may be needed to put these assignments before the above line:

```
\catcode \{=1  
\catcode \}=2
```

More fine-grained primitives control is possible and you can look up the details in section 10.3.15. For simplicity’s sake, this manual assumes that you have executed the `\directlua` command as given above.

The startup behaviour documented above is considered stable in the sense that there will not be backward-incompatible changes any more. We have promoted some rather generic P^DF_T_EX primitives to core L^AT_EX ones, and a few that we inherited from ALEPH (OMEGA) are also promoted. Effectively this means that we now only have the `tex`, `etex` and `luatex` sets left.

5.1.2 Experiments

There are a few extensions to the engine regarding the macro machinery. Some are already well tested but others are (still) experimental. Although they are likely to stay, their exact behaviour might evolve. Because L^AMETAT_EX is also used for experiments, this is not a problem. We can always decide to also add some of what is discussed here to L^AT_EX, but it will happen with a delay.

¹ At some point this function might be changed to return the whole list always



There are all kind of small improvements that might find their way into stock L^AT_EX: a few more helpers, some cleanup of code, etc. We'll see. In any case, if you play with these before they are declared stable, unexpected side effects are what you have to accept.

5.1.3 Version information

5.1.3.1 `\luatexbanner`, `\luatexversion` and `\luatexrevision`

There are three primitives to test the version of L^AT_EX (and L^AM^ET_EX):

PRIMITIVE	VALUE	EXPLANATION
<code>\luatexbanner</code>	This is LuaMetaTeX, Version 2.03.6	the banner reported on the command line
<code>\luatexversion</code>	203	a combination of major and minor number
<code>\luatexrevision</code>	6	the revision number, the current value is

A version is defined as follows:

- ▶ The major version is the integer result of `\luatexversion` divided by 100. The primitive is an 'internal variable', so you may need to prefix its use with `\the` depending on the context.
- ▶ The minor version is the two-digit result of `\luatexversion` modulo 100.
- ▶ The revision is reported by `\luatexrevision`. This primitive expands to a positive integer.
- ▶ The full version number consists of the major version, minor version and revision, separated by dots.

The L^AM^ET_EX version number starts at 2 in order to prevent a clash with L^AT_EX, and the version commands are the same. This is a way to indicate that these projects are related.

5.1.3.2 `\formatname`

The `\formatname` syntax is identical to `\jobname`. In I^NI^TE_X, the expansion is empty. Otherwise, the expansion is the value that `\jobname` had during the I^NI^TE_X run that dumped the currently loaded format. You can use this token list to provide your own version info.

5.2 UNICODE text support

5.2.1 Extended ranges

Text input and output is now considered to be UNICODE text, so input characters can use the full range of UNICODE ($2^{20} + 2^{16} - 1 = 0x10FFFF$). Later chapters will talk of characters and glyphs. Although these are not interchangeable, they are closely related. During typesetting, a character is always converted to a suitable graphic representation of that character in a specific font. However, while processing a list of to-be-typeset nodes, its contents may still be seen as



a character. Inside the engine there is no clear separation between the two concepts. Because the subtype of a glyph node can be changed in LUA it is up to the user. Subtypes larger than 255 indicate that font processing has happened.

A few primitives are affected by this, all in a similar fashion: each of them has to accommodate for a larger range of acceptable numbers. For instance, `\char` now accepts values between 0 and 1,114,111. This should not be a problem for well-behaved input files, but it could create incompatibilities for input that would have generated an error when processed by older \TeX -based engines. The affected commands with an altered initial (left of the equal sign) or secondary (right of the equal sign) value are: `\char`, `\lccode`, `\uccode`, `\hjcode`, `\catcode`, `\sfcode`, `\efcode`, `\lpcode`, `\rpcode`, `\chardef`.

As far as the core engine is concerned, all input and output to text files is UTF-8 encoded. Input files can be pre-processed using the reader callback. This will be explained in section ?? . Normalization of the UNICODE input is on purpose not built-in and can be handled by a macro package during callback processing. We have made some practical choices and the user has to live with those.

Output in byte-sized chunks can be achieved by using characters just outside of the valid UNICODE range, starting at the value 1,114,112 (0x110000). When the time comes to print a character $c \geq 1,114,112$, \LaTeX will actually print the single byte corresponding to c minus 1,114,112.

Contrary to other \TeX engines, the output to the terminal is as-is so there is no escaping with `^^`. We operate in a UTF universe.

5.2.2 `\Uchar`

The expandable command `\Uchar` reads a number between 0 and 1,114,111 and expands to the associated Unicode character.

5.2.3 Extended tables

All traditional \TeX and $\epsilon\text{-}\text{\TeX}$ registers can be 16-bit numbers. The affected commands are:

<code>\count</code>	<code>\countdef</code>	<code>\box</code>	<code>\wd</code>
<code>\dimen</code>	<code>\dimendef</code>	<code>\unhbox</code>	<code>\ht</code>
<code>\skip</code>	<code>\skipdef</code>	<code>\unvbox</code>	<code>\dp</code>
<code>\muskip</code>	<code>\muskipdef</code>	<code>\copy</code>	<code>\setbox</code>
<code>\marks</code>	<code>\toksdef</code>	<code>\unhcopy</code>	<code>\vsplit</code>
<code>\toks</code>	<code>\insert</code>	<code>\unvcopy</code>	

Fonts are loaded via LUA and a minimal amount of information is kept at the \TeX end. Sharing resources is up to the loaders. The engine doesn't really care about what a character (or glyph) number represents (an UNICODE or index) as it only is interested in dimensions.



5.3 Attributes

5.3.1 Nodes

When \TeX reads input it will interpret the stream according to the properties of the characters. Some signal a macro name and trigger expansion, others open and close groups, trigger math mode, etc. What's left over becomes the typeset text. Internally we get a linked list of nodes. Characters become glyph nodes that have for instance a font and char property and \kern 10pt becomes a kern node with a width property. Spaces are alien to \TeX as they are turned into glue nodes. So, a simple paragraph is mostly a mix of sequences of glyph nodes (words) and glue nodes (spaces). A node can have a subtype so that it can be recognized as for instance a space related glue.

The sequences of characters at some point are extended with disc nodes that relate to hyphenation. After that font logic can be applied and we get a list where some characters can be replaced, for instance multiple characters can become one ligature, and font kerns can be injected. This is driven by the font properties.

Boxes (like \hbox and \vbox) become hlist or vlist nodes with width, height, depth and shift properties and a pointer list to its actual content. Boxes can be constructed explicitly or can be the result of subprocesses. For instance, when lines are broken into paragraphs, the lines are a linked list of hlist nodes, possibly with glue and penalties in between.

Internally nodes have a number. This number is actually an index in the memory used to store nodes.

So, to summarize: all that you enter as content eventually becomes a node, often as part of a (nested) list structure. They have a relative small memory footprint and carry only the minimal amount of information needed. In traditional \TeX a character node only held the font and slot number, in $\text{LUA}\text{\TeX}$ we also store some language related information, the expansion factor, etc. Now that we have access to these nodes from LUA it makes sense to be able to carry more information with an node and this is where attributes kick in.

5.3.2 Attribute registers

Attributes are a completely new concept in $\text{LUA}\text{\TeX}$. Syntactically, they behave a lot like counters: attributes obey \TeX 's nesting stack and can be used after \the etc. just like the normal \count registers.

```
\attribute <16-bit number> <optional equals> <32-bit number>
\attributedef <csname> <optional equals> <16-bit number>
```

Conceptually, an attribute is either 'set' or 'unset'. Unset attributes have a special negative value to indicate that they are unset, that value is the lowest legal value: $-0x7FFFFFFF$ in hexadecimal, a.k.a. -2147483647 in decimal. It follows that the value $-0x7FFFFFFF$ cannot be used as a legal attribute value, but you *can* assign $-0x7FFFFFFF$ to 'unset' an attribute. All attributes start out in this 'unset' state in $\text{ini}\text{\TeX}$.

Attributes can be used as extra counter values, but their usefulness comes mostly from the fact that the numbers and values of all 'set' attributes are attached to all nodes created in their



scope. These can then be queried from any LUA code that deals with node processing. Further information about how to use attributes for node list processing from LUA is given in chapter 8.

Attributes are stored in a sorted (sparse) linked list that are shared when possible. This permits efficient testing and updating. You can define many thousands of attributes but normally such a large number makes no sense and is also not that efficient because each node carries a (possibly shared) link to a list of currently set attributes. But they are a convenient extension and one of the first extensions we implemented in L^AT_EX.

In L^AMETAT_EX we try to minimize the memory footprint and creation of these attribute lists more aggressive sharing them. This feature is still somewhat experimental.

5.3.3 Box attributes

Nodes typically receive the list of attributes that is in effect when they are created. This moment can be quite asynchronous. For example: in paragraph building, the individual line boxes are created after the `\par` command has been processed, so they will receive the list of attributes that is in effect then, not the attributes that were in effect in, say, the first or third line of the paragraph.

Similar situations happen in L^AT_EX regularly. A few of the more obvious problematic cases are dealt with: the attributes for nodes that are created during hyphenation, kerning and ligaturing borrow their attributes from their surrounding glyphs, and it is possible to influence box attributes directly.

When you assemble a box in a register, the attributes of the nodes contained in the box are unchanged when such a box is placed, unboxed, or copied. In this respect attributes act the same as characters that have been converted to references to glyphs in fonts. For instance, when you use attributes to implement color support, each node carries information about its eventual color. In that case, unless you implement mechanisms that deal with it, applying a color to already boxed material will have no effect. Keep in mind that this incompatibility is mostly due to the fact that separate specials and literals are a more unnatural approach to colors than attributes.

It is possible to fine-tune the list of attributes that are applied to a hbox, vbox or vtop by the use of the keyword `attr`. The `attr` keyword(s) should come before a `to` or `spread`, if that is also specified. An example is:

```
\attribute997=123
\attribute998=456
\setbox0=\hbox {Hello}
\setbox2=\hbox attr 999 = 789 attr 998 = -"7FFFFFFF{Hello}
```

Box 0 now has attributes 997 and 998 set while box 2 has attributes 997 and 999 set while the nodes inside that box will all have attributes 997 and 998 set. Assigning the maximum negative value causes an attribute to be ignored.

To give you an idea of what this means at the LUA end, take the following code:

```
for b=0,2,2 do
  for a=997, 999 do
```



```

    tex.sprint("box ", b, " : attr ",a," : ",tostring(tex.box[b]      [a]))
    tex.sprint("\\quad\\quad")
    tex.sprint("list ",b, " : attr ",a," : ",tostring(tex.box[b].list[a]))
    tex.sprint("\\par")
end
end

```

Later we will see that you can access properties of a node. The boxes here are so called `hlist` nodes that have a field `list` that points to the content. Because the attributes are a list themselves you can access them by indexing the node (here we do that with `[a]`). Running this snippet gives:

```

box 0 : attr 997 : 123    list 0 : attr 997 : 123
box 0 : attr 998 : 456    list 0 : attr 998 : 456
box 0 : attr 999 : nil    list 0 : attr 999 : nil
box 2 : attr 997 : 123    list 2 : attr 997 : 123
box 2 : attr 998 : nil    list 2 : attr 998 : 456
box 2 : attr 999 : 789    list 2 : attr 999 : nil

```

Because some values are not set we need to apply the `tostring` function here so that we get the word `nil`.

5.4 LUA related primitives

5.4.1 `\directlua`

In order to merge LUA code with \TeX input, a few new primitives are needed. The primitive `\directlua` is used to execute LUA code immediately. The syntax is

```
\directlua <general text>
```

The `<general text>` is expanded fully, and then fed into the LUA interpreter. After reading and expansion has been applied to the `<general text>`, the resulting token list is converted to a string as if it was displayed using `\the\toks`. On the LUA side, each `\directlua` block is treated as a separate chunk. In such a chunk you can use the `local` directive to keep your variables from interfering with those used by the macro package.

The conversion to and from a token list means that you normally can not use LUA line comments (starting with `--`) within the argument. As there typically will be only one ‘line’ the first line comment will run on until the end of the input. You will either need to use \TeX -style line comments (starting with `%`), or change the \TeX category codes locally. Another possibility is to say:

```

\begingroup
\endlinechar=10
\directlua ...
\endgroup

```

Then LUA line comments can be used, since \TeX does not replace line endings with spaces. Of course such an approach depends on the macro package that you use.



The `\directlua` command is expandable. Since it passes LUA code to the LUA interpreter its expansion from the \TeX viewpoint is usually empty. However, there are some LUA functions that produce material to be read by \TeX , the so called print functions. The most simple use of these is `tex.print(<string> s)`. The characters of the string `s` will be placed on the \TeX input buffer, that is, ‘before \TeX ’s eyes’ to be read by \TeX immediately. For example:

```
\count10=20
a\directlua{tex.print(tex.count[10]+5)}b
```

expands to

a25b

Here is another example:

```
$\pi = \directlua{tex.print(math.pi)}$
```

will result in

$\pi = 3.1415926535898$

Note that the expansion of `\directlua` is a sequence of characters, not of tokens, contrary to all \TeX commands. So formally speaking its expansion is null, but it places material on a pseudo-file to be immediately read by \TeX , as ε - \TeX ’s `\scantokens`. For a description of print functions look at section 10.3.13.

Because the `<general text>` is a chunk, the normal LUA error handling is triggered if there is a problem in the included code. The LUA error messages should be clear enough, but the contextual information is still pretty bad. Often, you will only see the line number of the right brace at the end of the code.

While on the subject of errors: some of the things you can do inside LUA code can break up `LUAMETATEX` pretty bad. If you are not careful while working with the node list interface, you may even end up with assertion errors from within the \TeX portion of the executable.

5.4.2 `\luaescapestring`

This primitive converts a \TeX token sequence so that it can be safely used as the contents of a LUA string: embedded backslashes, double and single quotes, and newlines and carriage returns are escaped. This is done by prepending an extra token consisting of a backslash with category code 12, and for the line endings, converting them to `n` and `r` respectively. The token sequence is fully expanded.

```
\luaescapestring <general text>
```

Most often, this command is not actually the best way to deal with the differences between \TeX and LUA. In very short bits of LUA code it is often not needed, and for longer stretches of LUA code it is easier to keep the code in a separate file and load it using LUA’s `dofile`:

```
\directlua { dofile("mysetups.lua") }
```

5.4.3 `\luafunction`, `\luafunctioncall` and `\luadef`

The `\directlua` commands involves tokenization of its argument (after picking up an optional name or number specification). The tokenlist is then converted into a string and given to LUA to



turn into a function that is called. The overhead is rather small but when you have millions of calls it can have some impact. For this reason there is a variant call available: `\luafunction`. This command is used as follows:

```
\directlua {
    local t = lua.get_functions_table()
    t[1] = function() tex.print("!") end
    t[2] = function() tex.print("?") end
}

\luafunction1
\luafunction2
```

Of course the functions can also be defined in a separate file. There is no limit on the number of functions apart from normal LUA limitations. Of course there is the limitation of no arguments but that would involve parsing and thereby give no gain. The function, when called in fact gets one argument, being the index, so in the following example the number 8 gets typeset.

```
\directlua {
    local t = lua.get_functions_table()
    t[8] = function(slot) tex.print(slot) end
}
```

The `\luafunctioncall` primitive does the same but is unexpandable, for instance in an `\edef`. In addition L^AT_EX provides a definer:

```

\luadef\MyFunctionA 1
\global\luadef\MyFunctionB 2
\protected\global\luadef\MyFunctionC 3
```

You should really use these commands with care. Some references get stored in tokens and assume that the function is available when that token expands. On the other hand, as we have tested this functionality in relative complex situations normal usage should not give problems.

5.4.4 `\luabytocode` and `\luabytodecall`

Analogue to the function callers discussed in the previous section we have byte code callers. Again the call variant is unexpandable.

```
\directlua {
    lua.bytecode[9998] = function(s)
        tex.sprint(s*token.scan_int())
    end
    lua.bytecode[5555] = function(s)
        tex.sprint(s*token.scan_dimen())
    end
}
```

This works with:




```
\luabytecode 9998 5 \luabytecode 5555 5sp
\luabytecodecall9998 5 \luabytecodecall5555 5sp
```

The variable `s` in the code is the number of the byte code register that can be used for diagnostic purposes. The advantage of bytecode registers over function calls is that they are stored in the format (but without upvalues).

5.5 Catcode tables

5.5.1 Catcodes

Catcode tables are a new feature that allows you to switch to a predefined catcode regime in a single statement. You can have lots of different tables, but if you need a dozen you might wonder what you're doing. . This subsystem is backward compatible: if you never use the following commands, your document will not notice any difference in behaviour compared to traditional \TeX . The contents of each catcode table is independent from any other catcode table, and its contents is stored and retrieved from the format file.

5.5.2 `\catcodetable`

`\catcodetable` <15-bit number>

The primitive `\catcodetable` switches to a different catcode table. Such a table has to be previously created using one of the two primitives below, or it has to be zero. Table zero is initialized by $\text{INIT}\TeX$.

5.5.3 `\initcatcodetable`

`\initcatcodetable` <15-bit number>

The primitive `\initcatcodetable` creates a new table with catcodes identical to those defined by $\text{INIT}\TeX$. The new catcode table is allocated globally: it will not go away after the current group has ended. If the supplied number is identical to the currently active table, an error is raised. The initial values are:

CATCODE	CHARACTER	EQUIVALENT	CATEGORY
0	\		escape
5	^^M	return	car_ret
9	^^@	null	ignore
10	<space>	space	spacer
11	a - z		letter
11	A - Z		letter
12	everything else		other
14	%		comment
15	^^?	delete	invalid_char



5.5.4 `\savecatcodetable`

`\savecatcodetable` <15-bit number>

`\savecatcodetable` copies the current set of catcodes to a new table with the requested number. The definitions in this new table are all treated as if they were made in the outermost level.

The new table is allocated globally: it will not go away after the current group has ended. If the supplied number is the currently active table, an error is raised.

5.6 Tokens, commands and strings

5.6.1 `\scantextokens`

The syntax of `\scantextokens` is identical to `\scantokens`. This primitive is a slightly adapted version of ε -TeX's `\scantokens`. The differences are:

- ▶ The last (and usually only) line does not have a `\endlinechar` appended.
- ▶ `\scantextokens` never raises an EOF error, and it does not execute `\everyeof` tokens.
- ▶ There are no ‘... while end of file ...’ error tests executed. This allows the expansion to end on a different grouping level or while a conditional is still incomplete.

5.6.2 `\toksapp`, `\tokspre`, `\etoksapp`, `\etokspre`, `\gtoksapp`, `\gtokspre`, `\xtoksapp`, `\xtokspre`

Instead of:

```
\toks0\expandafter{\the\toks0 foo}
```

you can use:

```
\etoksapp0{foo}
```

The pre variants prepend instead of append, and the e variants expand the passed general text. The g and x variants are global.

5.6.3 `\csstring`, `\beginsname` and `\lastnamedcs`

These are somewhat special. The `\csstring` primitive is like `\string` but it omits the leading escape character. This can be somewhat more efficient than stripping it afterwards.

The `\beginsname` primitive is like `\csname` but doesn't create a relaxed equivalent when there is no such name. It is equivalent to

```
\ifcsname foo\endcsname  
  \csname foo\endcsname  
\fi
```



The advantage is that it saves a lookup (don't expect much speedup) but more important is that it avoids using the `\if` test. The `\lastnamedcs` is one that should be used with care. The above example could be written as:

```
\ifcsname foo\endcsname
  \lastnamedcs
\fi
```

This is slightly more efficient than constructing the string twice (deep down in L^AT_EX this also involves some UTF8 juggling), but probably more relevant is that it saves a few tokens and can make code a bit more readable.

5.6.4 `\clearmarks`

This primitive complements the ϵ -T_EX mark primitives and clears a mark class completely, resetting all three connected mark texts to empty. It is an immediate command.

```
\clearmarks <16-bit number>
```

5.6.5 `\alignmark` and `\aligntab`

The primitive `\alignmark` duplicates the functionality of `#` inside alignment preambles, while `\aligntab` duplicates the functionality of `&`.

5.6.6 `\letcharcode`

This primitive can be used to assign a meaning to an active character, as in:

```
\def\foo{bar} \letcharcode123=\foo
```

This can be a bit nicer than using the uppercase tricks (using the property of `\uppercase` that it treats active characters special).

5.6.7 `\glet`

This primitive is similar to:

```
\protected\def\glet{\global\let}
```

but faster (only measurable with millions of calls) and probably more convenient (after all we also have `\gdef`).

5.6.8 `\expanded`, `\immediateassignment` and `\immediateassigned`

The `\expanded` primitive takes a token list and expands its content which can come in handy: it avoids a tricky mix of `\expandafter` and `\noexpand`. You can compare it with what happens inside the body of an `\edef`. But this kind of expansion still doesn't expand some primitive operations.



```
\newcount\NumberOfCalls
```

```
\def\TestMe{\advance\NumberOfCalls1 }
```

```
\edef\Tested{\TestMe foo:\the\NumberOfCalls}
```

```
\edef\Tested{\TestMe foo:\the\NumberOfCalls}
```

```
\edef\Tested{\TestMe foo:\the\NumberOfCalls}
```

\meaning\Tested

The result is a macro that has the not expanded code in its body:

```
macro:->\advance \NumberOfCalls 1 foo:0
```

Instead we can define \TestMe in a way that expands the assignment immediately. You need of course to be aware of preventing look ahead interference by using a space or \relax (often an expression works better as it doesn't leave an \relax).

```
\def\TestMe{\immediateassignment\advance\NumberOfCalls1 }
```

```
\edef\Tested{\TestMe foo:\the\NumberOfCalls}
```

```
\edef\Tested{\TestMe foo:\the\NumberOfCalls}
```

```
\edef\Tested{\TestMe foo:\the\NumberOfCalls}
```

\meaning\Tested

This time the counter gets updates and we don't see interference in the resulting \Tested macro:

```
macro:->foo:3
```

Here is a somewhat silly example of expanded comparison:

```
\def\expandeddoifelse#1#2#3#4%  
  {\immediateassignment\edef\tempa{#1}%  
   \immediateassignment\edef\tempb{#2}%  
   \ifx\tempa\tempb  
     \immediateassignment\def\next{#3}%  
   \else  
     \immediateassignment\def\next{#4}%  
   \fi  
  \next}
```

```
\edef\Tested  
  {(\expandeddoifelse{abc}{def}{yes}{nop}/%  
   \expandeddoifelse{abc}{abc}{yes}{nop})}
```

\meaning\Tested

It gives:

```
macro:->(nop/yes)
```



A variant is:

```
\def\expandeddoifelse#1#2#3#4%
{\immediateassigned{
  \edef\tempa{#1}%
  \edef\tempb{#2}%
}%
\ifx\tempa\tempb
  \immediateassignment\def\next{#3}%
\else
  \immediateassignment\def\next{#4}%
\fi
\next}
```

The possible error messages are the same as using assignments in preambles of alignments and after the `\accent` command. The supported assignments are the so called prefixed commands (except box assignments).

5.6.9 `\ignorepars`

This primitive is like `\ignorespaces` but also skips paragraph ending commands (normally `\par` and empty lines).

5.6.10 `\futureexpand`, `\futureexpandis`, `\futureexpandisap`

These commands are used as:

```
\futureexpand\sometoken\whenfound\whennotfound
```

When there is no match and a space was gobbled a space will be put back. The `is` variant doesn't do that while the `isap` even skips `\pars`. These characters stand for 'ignorespaces' and 'ignorespacesandpars'.

5.6.11 `\aftergrouped`

There is a new experimental feature that can inject multiple tokens to after the group ends. An example demonstrates its use:

```
{
  \aftergroup A \aftergroup B \aftergroup C
test 1 : }

{
  \aftergrouped{What comes next 1}
  \aftergrouped{What comes next 2}
  \aftergrouped{What comes next 3}
test 2 : }
```



```
{
  \aftergroup A \aftergrouped{What comes next 1}
  \aftergroup B \aftergrouped{What comes next 2}
  \aftergroup C \aftergrouped{What comes next 3}
test 3 : }
```

```
{
  \aftergrouped{What comes next 1} \aftergroup A
  \aftergrouped{What comes next 2} \aftergroup B
  \aftergrouped{What comes next 3} \aftergroup C
test 4 : }
```

This gives:

```
test 1 : ABC
test 2 : What comes next 1What comes next 2What comes next 3
test 3 : AWhat comes next 1BWhat comes next 2CWhat comes next 3
test 4 : What comes next 1AWhat comes next 2BWhat comes next 3C
```

5.7 Conditions

5.7.1 `\ifabsnum` and `\ifabsdim`

There are two tests that we took from PDF_T_EX:

```
\ifabsnum -10 = 10
  the same number
\fi
\ifabsdim -10pt = 10pt
  the same dimension
\fi
```

This gives

```
the same number the same dimension
```

5.7.2 `\ifcmpnum`, `\ifcmpdim`, `\ifnumval`, `\ifdimval`, `\ifchknum` and `\ifchkdim`

New are the ones that compare two numbers or dimensions:

```
\ifcmpnum 5 8 less \or equal \else more \fi
\ifcmpnum 5 5 less \or equal \else more \fi
\ifcmpnum 8 5 less \or equal \else more \fi
```

```
less equal more
```

and



```

\ifcmpdim 5pt 8pt less \or equal \else more \fi
\ifcmpdim 5pt 5pt less \or equal \else more \fi
\ifcmpdim 8pt 5pt less \or equal \else more \fi

```

less equal more

There are also some number and dimension tests. All four expose the `\else` branch when there is an error, but two also report if the number is less, equal or more than zero.

```

\ifnumval -123 \or < \or = \or > \or ! \else ? \fi
\ifnumval 0 \or < \or = \or > \or ! \else ? \fi
\ifnumval 123 \or < \or = \or > \or ! \else ? \fi
\ifnumval abc \or < \or = \or > \or ! \else ? \fi

```

```

\ifdimval -123pt \or < \or = \or > \or ! \else ? \fi
\ifdimval 0pt \or < \or = \or > \or ! \else ? \fi
\ifdimval 123pt \or < \or = \or > \or ! \else ? \fi
\ifdimval abcpt \or < \or = \or > \or ! \else ? \fi

```

< = > !

< = > !

```

\ifchknum -123 \or okay \else bad \fi
\ifchknum 0 \or okay \else bad \fi
\ifchknum 123 \or okay \else bad \fi
\ifchknum abc \or okay \else bad \fi

```

```

\ifchkdim -123pt \or okay \else bad \fi
\ifchkdim 0pt \or okay \else bad \fi
\ifchkdim 123pt \or okay \else bad \fi
\ifchkdim abcpt \or okay \else bad \fi

```

okay okay okay bad

okay okay okay bad

5.7.3 `\iftok` and `\ifcstok`

Comparing tokens and macros can be done with `\ifx`. Two extra test are provided in `LUAMETATEX`:

```

\def\ABC{abc} \def\DEF{def} \def\PQR{abc} \newtoks\XYZ \XYZ {abc}

```

```

\iftok{abc}{def}\relax (same) \else [different] \fi
\iftok{abc}{abc}\relax [same] \else (different) \fi
\iftok\XYZ {abc}\relax [same] \else (different) \fi

```

```

\ifcstok\ABC \DEF\relax (same) \else [different] \fi

```



```
\ifcstok\ABC \PQR\relax [same] \else (different) \fi
\ifcstok{abc}\ABC\relax [same] \else (different) \fi
```

```
[different][same][same]
[different][same][same]
```

You can check if a macro is defined as protected with `\ifprotected` while frozen macros can be tested with `\iffrozen`. A provisional `\ifusercmd` tests will check if a command is defined at the user level (and this one might evolve).

5.7.4 `\ifcondition`

This is a somewhat special one. When you write macros conditions need to be properly balanced in order to let \TeX 's fast branch skipping work well. This new primitive is basically a no-op flagged as a condition so that the scanner can recognize it as an if-test. However, when a real test takes place the work is done by what follows, in the next example `\something`.

```
\unexpanded\def\something#1#2%
  {\edef\tempa{#1}%
   \edef\tempb{#2}
   \ifx\tempa\tempb}

\ifcondition\something{a}{b}%
  \ifcondition\something{a}{a}%
    true 1
  \else
    false 1
  \fi
\else
  \ifcondition\something{a}{a}%
    true 2
  \else
    false 2
  \fi
\fi
```

If you are familiar with METAPOST, this is a bit like `vardef` where the macro has a return value. Here the return value is a test.

Experiments with something `\ifdef` actually worked ok but were rejected because in the end it gave no advantage so this generic one has to do. The `\ifcondition` test is basically is a no-op except when branches are skipped. However, when a test is expected, the scanner gobbles it and the next test result is used. Here is an other example:

```
\def\mytest#1%
  {\ifabsdim#1>0pt\else
   \expandafter \unless
   \fi
   \iftrue}
```




```
\ifcondition\mytest{10pt}\relax non-zero \else zero \fi
\ifcondition\mytest {0pt}\relax non-zero \else zero \fi
```

non-zero zero

The last expansion in a macro like `\mytest` has to be a condition and here we use `\unless` to negate the result.

5.7.5 `\orelse`

Sometimes you have successive tests that, when laid out in the source lead to deep trees. The `\ifcase` test is an exception. Experiments with `\ifcase` worked out fine but eventually were rejected because we have many tests so it would add a lot. As `LUAMETATEX` permitted more experiments, eventually an alternative was cooked up, one that has some restrictions but is relative lightweight. It goes like this:

```
\ifnum\count0<10
  less
\orelse\ifnum\count0=10
  equal
\else
  more
\fi
```

The `\orelse` has to be followed by one of the if test commands, except `\ifcondition`, and there can be an `\unless` in front of such a command. These restrictions make it possible to stay in the current condition (read: at the same level). If you need something more complex, using `\orelse` is probably unwise anyway. In case you wonder about performance, there is a little more checking needed when skipping branches but that can be neglected. There is some gain due to staying at the same level but that is only measurable when you runs tens of millions of complex tests and in that case it is very likely to drown in the real action. It's a convenience mechanism, in the sense that it can make your code look a bit easier to follow.

There is an nice side effect of this mechanism. When you define:

```
\def\quitcondition{\orelse\iffalse}
```

you can do this:

```
\ifnum\count0<10
  less
\orelse\ifnum\count0=10
  equal
  \quitcondition
  indeed
\else
  more
\fi
```



Of course it is only useful at the right level, so you might end up with cases like

```
\ifnum\count0<10
  less
\orelse\ifnum\count0=10
  equal
  \ifnum\count2=30
    \expandafter\quitcondition
  \fi
  indeed
\else
  more
\fi
```

5.7.6 `\ifprotected`, `\frozen`, `\iffrozen` and `\ifusercmd`

These checkers deal with control sequences. You can check if a command is a protected one, that is, defined with the `\protected` prefix. A command is frozen when it has been defined with the `\frozen` prefix. Beware: only macros can be frozen. A user command is a command that is not part of the predefined set of commands. This is an experimental command.

5.8 Boxes, rules and leaders

5.8.1 `\outputbox`

This integer parameter allows you to alter the number of the box that will be used to store the page sent to the output routine. Its default value is 255, and the acceptable range is from 0 to 65535.

```
\outputbox = 12345
```

5.8.2 `\vpack`, `\hpack` and `\tpack`

These three primitives are like `\vbox`, `\hbox` and `\vtop` but don't apply the related callbacks.

5.8.3 `\vsplit`

The `\vsplit` primitive has to be followed by a specification of the required height. As alternative for the `to` keyword you can use `upto` to get a split of the given size but result has the natural dimensions then.

5.8.4 Images and reused box objects

In original \TeX image support is dealt with via specials. It's not a native feature of the engine. All that \TeX cares about is dimensions, so in practice that meant: using a box with known dimensions



that wraps a special that instructs the backend to include an image. The wrapping is needed because a special itself is a whatsit and as such has no dimensions.

In PDF_T_E_X a special whatsit for images was introduced and that one *has* dimensions. As a consequence, in several places where the engine to deal with the dimensions of nodes, it now has to check the details of whatsits. By inheriting code from PDF_T_E_X, the LUAT_E_X engine also had that property. However, at some point this approach was abandoned and a more natural trick was used: images (and box resources) became a special kind of rules, and as rules already have dimensions, the code could be simplified.

When direction nodes and localpar nodes also became first class nodes, whatsits again became just that: nodes representing whatever you want, but without dimensions, and therefore they could again be ignored when dimensions mattered. And, because images were disguised as rules, as mentioned ,their dimensions automatically were taken into account. This seperation between front and backend cleaned up the code base already quite a bit.

In LUAMETAT_E_X we still have the image specific subtypes for rules, but the engine never looks at subtypes of rules. That was up to the backend. This means that image support is not present in LUAMETAT_E_X. When an image specification was parsed the special properties, like the filename, or additional attributes, were stored in the backend and all that LUAT_E_X does is registering a reference to an image s specification in the rule node. But, having no backend means nothing is stored, which in turn would make the image inclusion primitives kind of weird.

Therefore you need to realize that contrary to LUAT_E_X, *in LUAMETAT_E_X support for images and box reuse is not built in!* However, we can assume that an implementation uses rules in a similar fashion as LUAT_E_X does. So, you can still consider images and box reuse to be core concepts. Here we just mention the primitives that LUAT_E_X provides. They are not available in the engine but can of course be implemented in LUA.

COMMAND	EXPLANATION
<code>\saveboxresource</code>	save the box as an object to be included later
<code>\saveimageresource</code>	save the image as an object to be included later
<code>\useboxresource</code>	include the saved box object here (by index)
<code>\useimageresource</code>	include the saved image object here (by index)
<code>\lastsavedboxresourceindex</code>	the index of the last saved box object
<code>\lastsavedimageresourceindex</code>	the index of the last saved image object
<code>\lastsavedimageresourcepages</code>	the number of pages in the last saved image object

An implementation probably should accepts the usual optional dimension parameters for `\use...resource` in the same format as for rules. With images, these dimensions are then used instead of the ones given to `\useimageresource` but the original dimensions are not overwritten, so that a `\useimageresource` without dimensions still provides the image with dimensions defined by `\saveimageresource`. These optional parameters are not implemented for `\saveboxresource`.

```
\useimageresource width 20mm height 10mm depth 5mm \lastsavedimageresourceindex
\useboxresource   width 20mm height 10mm depth 5mm \lastsavedboxresourceindex
```

Examples or optional entries are `attr` and `resources` that accept a token list, and the `type` key. When set to non-zero the `/Type` entry is omitted. A value of 1 or 3 still writes a `/BBox`, while 2 or



3 will write a `/Matrix`. But, as said: this is entirely up to the backend. Generic macro packages (like `tikz`) can use these assumed primitives so one can best provide them. It is probably, for historic reasons, the only more or less standardized image inclusion interface one can expect to work in all macro packages.

5.8.5 `\hpack`, `\vpack` and `\tpack`

These three primitives are the equivalents of `\hbox`, `\vbox` and `\vtop` but they don't trigger the packaging related callbacks. Of course one never know if content needs a treatment so using them should be done with care.

5.8.6 `\nohrule` and `\novrule`

Because introducing a new keyword can cause incompatibilities, two new primitives were introduced: `\nohrule` and `\novrule`. These can be used to reserve space. This is often more efficient than creating an empty box with fake dimensions.

5.8.7 `\gleaders`

This type of leaders is anchored to the origin of the box to be shipped out. So they are like normal `\leaders` in that they align nicely, except that the alignment is based on the *largest* enclosing box instead of the *smallest*. The `g` stresses this global nature.

5.9 Languages

5.9.1 `\hyphenationmin`

This primitive can be used to set the minimal word length, so setting it to a value of 5 means that only words of 6 characters and more will be hyphenated, of course within the constraints of the `\lefthyphenmin` and `\righthyphenmin` values (as stored in the glyph node). This primitive accepts a number and stores the value with the language.

5.9.2 `\boundary`, `\noboundary`, `\protrusionboundary` and `\wordboundary`

The `\noboundary` command is used to inject a `whatsit` node but now injects a normal node with type `boundary` and subtype 0. In addition you can say:

```
x\boundary 123\relax y
```

This has the same effect but the subtype is now 1 and the value 123 is stored. The traditional ligature builder still sees this as a cancel boundary directive but at the LUA end you can implement different behaviour. The added benefit of passing this value is a side effect of the generalization. The subtypes 2 and 3 are used to control protrusion and word boundaries in hyphenation and have related primitives.



5.10 Control and debugging

5.10.1 Tracing

If `\tracingonline` is larger than 2, the node list display will also print the node number of the nodes.

5.10.2 `\lastnodetype`, `\lastnodesubtype`, `\currentifttype` and `\internalcodesmode`.

The ε -TeX command `\lastnodetype` is limited to some nodes. When the parameter `\internalcodesmode` is set to a non-zero value the normal (internally used) numbers are reported. The same is true for `\currentifttype`, as we have more conditionals and also use a different order. The `\lastnodesubtype` is a bonus.

5.11 Files

5.11.1 File syntax

LUAMETATEX will accept a braced argument as a file name:

```
\input {plain}  
\openin 0 {plain}
```

This allows for embedded spaces, without the need for double quotes. Macro expansion takes place inside the argument.

The `\tracingfonts` primitive that has been inherited from PDFTEX has been adapted to support variants in reporting the font. The reason for this extension is that a csname not always makes sense. The zero case is the default.

VALUE	REPORTED
0	<code>\foo xyz</code>
1	<code>\foo (bar)</code>
2	<code><bar> xyz</code>
3	<code><bar @ ..pt> xyz</code>
4	<code><id></code>
5	<code><id: bar></code>
6	<code><id: bar @ ..pt> xyz</code>

5.11.2 Writing to file

You can now open upto 127 files with `\openout`. When no file is open writes will go to the console and log. The write related primitives have to be implemented as part of a backend! As



a consequence a system command is no longer possible but one can use `os.execute` to do the same.

5.12 Math

We will cover math extensions in its own chapter because not only the font subsystem and spacing model have been enhanced (thereby introducing many new primitives) but also because some more control has been added to existing functionality. Much of this relates to the different approaches of traditional \TeX fonts and OPENTYPE math.

5.13 Fonts

Like math, we will cover fonts extensions in its own chapter. Here we stick to mentioning that loading fonts is different in $\text{\text{LUAMETAT}\TeX}$. As in $\text{\text{LUAT}\TeX}$ we have the extra primitives `\fontid` and `\setfontid`, `\noligs` and `\nokerns`, and `\nospaces`. The other new primitives in $\text{\text{LUAT}\TeX}$ have been dropped.

5.14 Directions

5.14.1 Two directions

The directional model in $\text{\text{LUAMETAT}\TeX}$ is a simplified version the the model used in $\text{\text{LUAT}\TeX}$. In fact, not much is happening at all: we only register a change in direction.

5.14.2 How it works

The approach is that we try to make node lists balanced but also try to avoid some side effects. What happens is quite intuitive if we forget about spaces (turned into glue) but even there what happens makes sense if you look at it in detail. However that logic makes in-group switching kind of useless when no proper nested grouping is used: switching from right to left several times nested, results in spacing ending up after each other due to nested mirroring. Of course a sane macro package will manage this for the user but here we are discussing the low level injection of directional information.

This is what happens:

```
\textdirection 1 nur {\textdirection 0 run \textdirection 1 NUR} nur
```

This becomes stepwise:

```
injected: [push 1]nur {[push 0]run [push 1]NUR} nur
balanced: [push 1]nur {[push 0]run [pop 0][push 1]NUR[pop 1]} nur[pop 0]
result   : run {RUNrun } run
```

And this:



```
\textdirection 1 nur {nur \textdirection 0 run \textdirection 1 NUR} nur
```

becomes:

```
injected: [+TRT]nur {nur [+TLT]run [+TRT]NUR} nur
balanced: [+TRT]nur {nur [+TLT]run [-TLT][+TRT]NUR[-TRT]} nur[-TRT]
result   : run {run RUNrun } run
```

Now, in the following examples watch where we put the braces:

```
\textdirection 1 nur {\textdirection 0 run} {\textdirection 1 NUR}} nur
```

This becomes:

```
run RUN run run
```

Compare this to:

```
\textdirection 1 nur {\textdirection 0 run }{\textdirection 1 NUR}} nur
```

Which renders as:

```
run RUNrun run
```

So how do we deal with the next?

```
\def\ltr{\textdirection 0\relax}
\def\rtl{\textdirection 1\relax}

run {\rtl nur {\ltr run \rtl NUR \ltr run \rtl NUR} nur}
run {\ltr run {\rtl nur \ltr RUN \rtl nur \ltr RUN} run}
```

It gets typeset as:

```
run run RUNrun RUNrun run
run run runRUN runRUN run
```

We could define the two helpers to look back, pick up a skip, remove it and inject it after the dir node. But that way we loose the subtype information that for some applications can be handy to be kept as-is. This is why we now have a variant of `\textdirection` which injects the balanced node before the skip. Instead of the previous definition we can use:

```
\def\ltr{\linedirection 0\relax}
\def\rtl{\linedirection 1\relax}
```

and this time:

```
run {\rtl nur {\ltr run \rtl NUR \ltr run \rtl NUR} nur}
run {\ltr run {\rtl nur \ltr RUN \rtl nur \ltr RUN} run}
```

comes out as a properly spaced:

```
run run RUN run RUN run run
```



run run run RUN run RUN run

Anything more complex than this, like combination of skips and penalties, or kerns, should be handled in the input or macro package because there is no way we can predict the expected behaviour. In fact, the `\linedir` is just a convenience extra which could also have been implemented using node list parsing.

5.14.3 Controlling glue with `\breakafterdirmode`

Glue after a `dir` node is ignored in the linebreak decision but you can bypass that by setting `\breakafterdirmode` to 1. The following table shows the difference. Watch your spaces.

	0	1
<code>pre {\textdirection 0 xxx} post</code>	pre xxx post	pre xxx post
<code>pre {\textdirection 0 xxx }post</code>	pre xxx post	pre xxx post
<code>pre{ \textdirection 0 xxx} post</code>	pre xxx post	pre xxx post
<code>pre{ \textdirection 0 xxx }post</code>	pre xxx post	pre xxx post
<code>pre { \textdirection 0 xxx } post</code>	pre xxx post	pre xxx post
<code>pre {\textdirection 0\relax \space xxx} post</code>	pre xxx post	pre xxx post

5.14.4 Controlling parshapes with `\shapemode`

Another adaptation to the ALEPH directional model is control over shapes driven by `\hangindent` and `\parshape`. This is controlled by a new parameter `\shapemode`:

VALUE	\HANGINDENT	\PARSHAPE
0	normal	normal
1	mirrored	normal
2	normal	mirrored
3	mirrored	mirrored

The value is reset to zero (like `\hangindent` and `\parshape`) after the paragraph is done with. You can use negative values to prevent this. In figure 5.1 a few examples are given.



We thrive in information-thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsise, winnow the wheat from the chaff and separate the sheep from the goats.

TLT: hangindent

ruo fo esuaceb sdlrow kciht-noitamrofni ni evirht eW
 -nis ,tide ,tceles ot yticapac yadyreve dna suolevram
 -rah ,egrem ,riap ,puorg ,thgilhgih ,erutcurts ,tuo elg
 ,nwod liob ,ecuder ,esnednoc ,ezinagro ,sucof ,ezisehtnys ,ezinom
 ,otni kool ,nacs ,tcartsba ,tsil ,yfissalc ,golatac ,ezirogetac ,esoohc
 ,kciip ,elohnoegip ,neercs ,hsiugnitsid ,etanimircsid ,etalosi ,ezilaedi
 ,knuhc ,htooms ,piks ,pmul ,retlfi ,tcepsni ,dnelb ,etargetni ,tros ,revo
 ,meti ,ezirammus ,eniltuo ,etagergga ,retsulc ,etamixorppa ,egareva
 ,hguorht fael ,otni ecnalg ,esworb ,hguorht pifl ,otni pid ,weiver ,ezi
 morf taehw eht wonniw ,ezisponys ,naelg ,etaremmune ,enfier ,miks
 ,staog eht morf peehs eht etarapes dna ffahc eht.

TRT: hangindent mode 0

ruo fo esuaceb sdlrow kciht-noitamrofni ni evirht eW
 -nis ,tide ,tceles ot yticapac yadyreve dna suolevram
 -rah ,egrem ,riap ,puorg ,thgilhgih ,erutcurts ,tuo elg
 ,nwod liob ,ecuder ,esnednoc ,ezinagro ,sucof ,ezisehtnys ,ezinom
 ,otni kool ,nacs ,tcartsba ,tsil ,yfissalc ,golatac ,ezirogetac ,esoohc
 ,kciip ,elohnoegip ,neercs ,hsiugnitsid ,etanimircsid ,etalosi ,ezilaedi
 ,knuhc ,htooms ,piks ,pmul ,retlfi ,tcepsni ,dnelb ,etargetni ,tros ,revo
 ,meti ,ezirammus ,eniltuo ,etagergga ,retsulc ,etamixorppa ,egareva
 ,hguorht fael ,otni ecnalg ,esworb ,hguorht pifl ,otni pid ,weiver ,ezi
 morf taehw eht wonniw ,ezisponys ,naelg ,etaremmune ,enfier ,miks
 ,staog eht morf peehs eht etarapes dna ffahc eht.

TRT: hangindent mode 1 & 3

We thrive in information-thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsise, winnow the wheat from the chaff and separate the sheep from the goats.

TLT: parshape

-ram ruo fo esuaceb sdlrow kciht-noitamrofni ni evirht eW
 ,tuo elgnis ,tide ,tceles ot yticapac yadyreve dna suolev
 -nys ,ezinomrah ,egrem ,riap ,puorg ,thgilhgih ,erutcurts
 ,ogetac ,esoohc ,nwod liob ,ecuder ,esnednoc ,ezinagro ,sucof ,eziseht
 ,etalosi ,ezilaedi ,otni kool ,nacs ,tcartsba ,tsil ,yfissalc ,golatac ,ezir
 ,etargetni ,tros ,revo kciip ,elohnoegip ,neercs ,hsiugnitsid ,etanimircsid
 ,etamixorppa ,egareva ,knuhc ,htooms ,piks ,pmul ,retlfi ,tcepsni ,dnelb
 pifl ,otni pid ,weiver ,ezimeti ,ezirammus ,eniltuo ,etagergga ,retsulc
 ,etaremmune ,enfier ,miks ,hguorht fael ,otni ecnalg ,esworb ,hguorht
 eht etarapes dna ffahc eht morf taehw eht wonniw ,ezisponys ,naelg
 ,staog eht morf peehs

TRT: parshape mode 0

-ram ruo fo esuaceb sdlrow kciht-noitamrofni ni evirht eW
 ,tuo elgnis ,tide ,tceles ot yticapac yadyreve dna suolev
 -nys ,ezinomrah ,egrem ,riap ,puorg ,thgilhgih ,erutcurts
 ,ogetac ,esoohc ,nwod liob ,ecuder ,esnednoc ,ezinagro ,sucof ,eziseht
 ,etalosi ,ezilaedi ,otni kool ,nacs ,tcartsba ,tsil ,yfissalc ,golatac ,ezir
 ,etargetni ,tros ,revo kciip ,elohnoegip ,neercs ,hsiugnitsid ,etanimircsid
 ,etamixorppa ,egareva ,knuhc ,htooms ,piks ,pmul ,retlfi ,tcepsni ,dnelb
 pifl ,otni pid ,weiver ,ezimeti ,ezirammus ,eniltuo ,etagergga ,retsulc
 ,etaremmune ,enfier ,miks ,hguorht fael ,otni ecnalg ,esworb ,hguorht
 eht etarapes dna ffahc eht morf taehw eht wonniw ,ezisponys ,naelg
 ,staog eht morf peehs

TRT: parshape mode 2 & 3

Figure 5.1 The effect of shapemode.

We have `\pardirection`, `\textdirection`, `\mathdirection` and `\linedirection` that is like `\textdirection` but with some additional (inline) glue checking.

5.14.5 Orientations

As mentioned, the difference with L^AT_EX is that we only have numeric directions and that there are only two: left-to-right (0) and right-to-left (1). The direction of a box is set with `direction`.

In addition to that boxes can now have an orientation keyword followed by optional `xoffset` and/or `yoffset` keywords. The offsets don't have consequences for the dimensions. The alternatives `xmove` and `ymove` on the contrary are reflected in the dimensions. Just play with them. The offsets and moves only are accepted when there is also an orientation, so no time is wasted on testing for these rarely used keywords. There are related primitives `\box...` that set these properties.

As these are experimental it will not be explained here (yet). They are covered in the descriptions of the development of L^AM^ET_AT_EX: articles and/or documents in the C^ON^TE^XT distribution. For now it is enough to know that the orientation can be up, down, left or right (rotated) and that it has some anchoring variants. Combined with the offsets this permits macro writers to provide solutions for top-down and bottom-up writing directions, something that is rather macro package



specific and used for scripts that need manipulations anyway. The ‘old’ vertical directions were never okay and therefore not used.

There are a couple of properties in boxes that you can set and query but that only really take effect when the backend supports them. When usage on CON_TE_XT shows that is’t okay, they will become official, so we just mention them: `\boxdirection`, `\boxattr`, `\boxorientation`, `\boxxoffset`, `\boxyoffset`, `\boxxmove`, `\boxymove` and `\boxtotal`.

5.15 Expressions

The `*expr` parsers now accept `:` as operator for integer division (the `/` operators does rounding. This can be used for division compatible with `\divide`. I’m still wondering if adding a couple of bit operators makes sense (for integers).

5.16 Nodes

The ε -T_EX primitive `\lastnodetype` is not honest in reporting the internal numbers as it uses its own values. But you can set `\internalcodesmode` to a non-zero value to get the real id’s instead. In addition there is `\lastnodesubtype`.

Another last one is `\lastnamedcs` which holds the last match but this one should be used with care because one never knows if in the meantime something else ‘last’ has been seen.



6 Fonts

6.1 Introduction

Only traditional font support is built in, anything more needs to be implemented in LUA. This is conform the L^AT_EX philosophy. When you pass a font to the frontend only the dimensions matter, as these are used in typesetting, and optionally ligatures and kerns when you rely on the built-in font handler. For math some extra data is needed, like information about extensibles and next in size glyphs. You can of course put more information in your LUA tables because when such a table is passed to T_EX only that what is needed is filtered from it.

Because there is no built-in backend, virtual font information is not used. If you want to be compatible you'd better make sure that your tables are okay, and in that case you can best consult the L^AT_EX manual. For instance, parameters like `extend` are backend related and the standard L^AT_EX backend sets the standard here.

6.2 Defining fonts

All T_EX fonts are represented to LUA code as tables, and internally as C structures. All keys in the table below are saved in the internal font structure if they are present in the table passed to `font.define`. When the callback is set, which is needed for `\font` to work, its function gets the name and size passed, and it has to return a valid font identifier (a positive number).

For the engine to work well, the following information has to be present at the font level:

KEY	VALUE TYPE	DESCRIPTION
<code>name</code>	string	metric (file) name
<code>characters</code>	table	the defined glyphs of this font
<code>designsize</code>	number	expected size (default: 655360 == 10pt)
<code>fonts</code>	table	locally used fonts
<code>hyphenchar</code>	number	default: T _E X's <code>\hyphenchar</code>
<code>parameters</code>	hash	default: 7 parameters, all zero
<code>size</code>	number	the required scaling (by default the same as <code>designsize</code>)
<code>skewchar</code>	number	default: T _E X's <code>\skewchar</code>
<code>stretch</code>	number	the 'stretch'
<code>shrink</code>	number	the 'shrink'
<code>step</code>	number	the 'step'
<code>nomath</code>	boolean	this key allows a minor speedup for text fonts; if it is present and true, then L ^A T _E X will not check the character entries for math-specific keys
<code>oldmath</code>	boolean	this key flags a font as representing an old school T _E X math font and disables the OPENTYPE code path

The `parameters` is a hash with mixed key types. There are seven possible string keys, as well as a number of integer indices (these start from 8 up). The seven strings are actually used instead of the bottom seven indices, because that gives a nicer user interface.



The names and their internal remapping are:

NAME	REMAPPING
slant	1
space	2
space_stretch	3
space_shrink	4
x_height	5
quad	6
extra_space	7

The characters table is a list of character hashes indexed by an integer number. The number is the ‘internal code’ \TeX knows this character by. For proper paragraph building and math rendering the following fields can be present in an entry in the characters table. You can of course add all kind of extra fields. The engine only uses those that it needs for typesetting a paragraph or formula.

Each character hash itself is a hash. For example, here is the character ‘f’ (decimal 102) in the font `cmr10` at 10pt. The numbers that represent dimensions are in scaled points.

```
[102] = {
  ["width"] = 200250,
  ["height"] = 455111,
  ["depth"] = 0,
  ["italic"] = 50973,
  ["kerns"] = {
    [63] = 50973,
    [93] = 50973,
    [39] = 50973,
    [33] = 50973,
    [41] = 50973
  },
  ["ligatures"] = {
    [102] = { ["char"] = 11, ["type"] = 0 },
    [108] = { ["char"] = 13, ["type"] = 0 },
    [105] = { ["char"] = 12, ["type"] = 0 }
  }
}
```

Providing ligatures and kerns this way permits \TeX to construct ligatures and add inter-character kerning. However, normally you will use an `OPENTYPE` font in combination with `LUA` code that does this. In `CON \TeX TEXT` we have base mode that uses the engine, and node mode that uses `LUA`. A monospaced font normally has no ligatures and kerns and is normally not processed at all.

KEY	TYPE	DESCRIPTION
width	number	width in sp (default 0)
height	number	height in sp (default 0)
depth	number	depth in sp (default 0)



<code>italic</code>	number	italic correction in sp (default 0)
<code>top_accent</code>	number	top accent alignment place in sp (default zero)
<code>bot_accent</code>	number	bottom accent alignment place, in sp (default zero)
<code>left_protruding</code>	number	left protruding factor (<code>\lpcode</code>)
<code>right_protruding</code>	number	right protruding factor (<code>\rpcode</code>)
<code>expansion_factor</code>	number	expansion factor (<code>\efcode</code>)
<code>next</code>	number	'next larger' character index
<code>extensible</code>	table	constituent parts of an extensible recipe
<code>vert_variants</code>	table	constituent parts of a vertical variant set
<code>horiz_variants</code>	table	constituent parts of a horizontal variant set
<code>kerns</code>	table	kerning information
<code>ligatures</code>	table	ligaturing information
<code>mathkern</code>	table	math cut-in specifications

Two very special string indexes can be used also: `left_boundary` is a virtual character whose ligatures and kerns are used to handle word boundary processing. `right_boundary` is similar but not actually used for anything (yet).

The values of `top_accent`, `bot_accent` and `mathkern` are used only for math accent and superscript placement, see page 65 in this manual for details. The values of `left_protruding` and `right_protruding` are used only when `\protrudechars` is non-zero. Whether or not `expansion_factor` is used depends on the font's global expansion settings, as well as on the value of `\adjustspacing`.

A math character can have a `next` field that points to a next larger shape. However, the presence of `extensible` will overrule `next`, if that is also present. The `extensible` field in turn can be overruled by `vert_variants`, the OPENTYPE version. The `extensible` table is very simple:

KEY	TYPE	DESCRIPTION
<code>top</code>	number	top character index
<code>mid</code>	number	middle character index
<code>bot</code>	number	bottom character index
<code>rep</code>	number	repeatable character index

The `horiz_variants` and `vert_variants` are arrays of components. Each of those components is itself a hash of up to five keys:

KEY	TYPE	EXPLANATION
<code>glyph</code>	number	The character index. Note that this is an encoding number, not a name.
<code>extender</code>	number	One (1) if this part is repeatable, zero (0) otherwise.
<code>start</code>	number	The maximum overlap at the starting side (in scaled points).
<code>end</code>	number	The maximum overlap at the ending side (in scaled points).
<code>advance</code>	number	The total advance width of this item. It can be zero or missing, then the natural size of the glyph for character component is used.

The `kerns` table is a hash indexed by character index (and 'character index' is defined as either a non-negative integer or the string value `right_boundary`), with the values of the kerning to be applied, in scaled points.



The ligatures table is a hash indexed by character index (and ‘character index’ is defined as either a non-negative integer or the string value `right_boundary`), with the values being yet another small hash, with two fields:

KEY	TYPE	DESCRIPTION
type	number	the type of this ligature command, default 0
char	number	the character index of the resultant ligature

The `char` field in a ligature is required. The `type` field inside a ligature is the numerical or string value of one of the eight possible ligature types supported by \TeX . When \TeX inserts a new ligature, it puts the new glyph in the middle of the left and right glyphs. The original left and right glyphs can optionally be retained, and when at least one of them is kept, it is also possible to move the new ‘insertion point’ forward one or two places. The glyph that ends up to the right of the insertion point will become the next ‘left’.

TEXTUAL (KNUTH)	NUMBER	STRING	RESULT
<code>l + r =: n</code>	0	<code>=:</code>	<code> n</code>
<code>l + r =: n</code>	1	<code>=: </code>	<code> nr</code>
<code>l + r =: n</code>	2	<code> =:</code>	<code> ln</code>
<code>l + r =: n</code>	3	<code> =: </code>	<code> lnr</code>
<code>l + r =: > n</code>	5	<code>=: ></code>	<code>n r</code>
<code>l + r =: > n</code>	6	<code> =: ></code>	<code>l n</code>
<code>l + r =: > n</code>	7	<code> =: ></code>	<code>l nr</code>
<code>l + r =: >> n</code>	11	<code> =: >></code>	<code>ln r</code>

The default value is 0, and can be left out. That signifies a ‘normal’ ligature where the ligature replaces both original glyphs. In this table the `|` indicates the final insertion point.

6.3 Virtual fonts

Virtual fonts have been introduced to overcome limitations of good old \TeX . They were mostly use for providing a direct mapping from for instance accented characters onto a glyph. The backend was responsible for turning a reference to a character slot into a real glyph, possibly constructed from other glyphs. In our case there is no backend so there is also no need to pass this information through \TeX . But it can of course be part of the font information and because it is a kind of standard, we describe it here.

A character is virtual when it has a `commands` array as part of the data. A virtual character can itself point to virtual characters but be careful with nesting as you can create loops and overflow the stack (which often indicates an error anyway).

At the font level there can be a `fonts` an (indexed) `LUA` table. The values are one- or two-key hashes themselves, each entry indicating one of the base fonts in a virtual font. In case your font is referring to itself in for instance a virtual font, you can use the `slot` command with a zero font reference, which indicates that the font itself is used. So, a table looks like this:

```
fonts = {
  { name = "ptmr8a", size = 655360 },
```



```

{ name = "psyr", size = 600000 },
{ id = 38 }
}

```

The first referenced font (at index 1) in this virtual font is ptmr8a loaded at 10pt, and the second is psyr loaded at a little over 9pt. The third one is a previously defined font that is known to L^AT_EX as font id 38. The array index numbers are used by the character command definitions that are part of each character.

The commands array is a hash where each item is another small array, with the first entry representing a command and the extra items being the parameters to that command. The allowed commands and their arguments are:

COMMAND	ARGUMENTS	TYPE	DESCRIPTION
font	1	number	select a new font from the local fonts table
char	1	number	typeset this character number from the current font, and move right by the character's width
node	1	node	output this node (list), and move right by the width of this list
slot	2	2 numbers	a shortcut for the combination of a font and char command
push	0		save current position
nop	0		do nothing
pop	0		pop position
rule	2	2 numbers	output a rule $ht * wd$, and move right.
down	1	number	move down on the page
right	1	number	move right on the page
special	1	string	output a \special command
pdf	2	2 strings	output a PDF literal, the first string is one of origin, page, text, font, direct or raw; if you have one string only origin is assumed
lua	1	string, function	execute a LUA script when the glyph is embedded; in case of a function it gets the font id and character code passed
image	1	image	output an image (the argument can be either an <image> variable or an image_spec table)
comment	any	any	the arguments of this command are ignored

When a font id is set to 0 then it will be replaced by the currently assigned font id. This prevents the need for hackery with future id's.

The pdf option also accepts a mode keyword in which case the third argument sets the mode. That option will change the mode in an efficient way (passing an empty string would result in an extra empty lines in the PDF file. This option only makes sense for virtual fonts. The font mode only makes sense in virtual fonts. Modes are somewhat fuzzy and partially inherited from PDF_T_EX.

MODE	DESCRIPTION
origin	enter page mode and set the position



page	enter page mode
text	enter text mode
font	enter font mode (kind of text mode, only in virtual fonts)
always	finish the current string and force a transform if needed
raw	finish the current string

You always need to check what PDF code is generated because there can be all kind of interferences with optimization in the backend and fonts are complicated anyway. Here is a rather elaborate glyph commands example using such keys:

```
...
commands = {
  { "push" },                -- remember where we are
  { "right", 5000 },         -- move right about 0.08pt
  { "font", 3 },            -- select the fonts[3] entry
  { "char", 97 },           -- place character 97 (ASCII 'a')
  -- { "slot", 2, 97 },      -- an alternative for the previous two
  { "pop" },               -- go all the way back
  { "down", -200000 },      -- move upwards by about 3pt
  { "special", "pdf: 1 0 0 rg" } -- switch to red color
  -- { "pdf", "origin", "1 0 0 rg" } -- switch to red color (alternative)
  { "rule", 500000, 20000 } -- draw a bar
  { "special", "pdf: 0 g" }  -- back to black
  -- { "pdf", "origin", "0 g" }  -- back to black (alternative)
}
...
```

The default value for font is always 1 at the start of the commands array. Therefore, if the virtual font is essentially only a re-encoding, then you do usually not have created an explicit 'font' command in the array.

Rules inside of commands arrays are built up using only two dimensions: they do not have depth. For correct vertical placement, an extra down command may be needed.

Regardless of the amount of movement you create within the commands, the output pointer will always move by exactly the width that was given in the width key of the character hash. Any movements that take place inside the commands array are ignored on the upper level.

The special can have a pdf:, pdf:origin:, pdf:page:, pdf:direct: or pdf:raw: prefix. When you have to concatenate strings using the pdf command might be more efficient.

The fields mentioned above can be found in external fonts. It is good to keep in mind that we can extend this model, given that the backend knows what to do with it.

6.4 Additional T_EX commands

6.4.1 Font syntax

LUAT_EX will accept a braced argument as a font name:




```
\font\myfont = {cmr10}
```

This allows for embedded spaces, without the need for double quotes. Macro expansion takes place inside the argument.

6.4.2 `\fontid` and `\setfontid`

```
\fontid\font
```

This primitive expands into a number. It is not a register so there is no need to prefix with `\number` (and using `\the` gives an error). The currently used font id is 29. Here are some more:

STYLE	COMMAND	FONT ID
normal	<code>\tf</code>	29
bold	<code>\bf</code>	38
italic	<code>\it</code>	59
bold italic	<code>\bi</code>	75

These numbers depend on the macro package used because each one has its own way of dealing with fonts. They can also differ per run, as they can depend on the order of loading fonts. For instance, when in `CONTEX`T virtual math `UNICODE` fonts are used, we can easily get over a hundred ids in use. Not all ids have to be bound to a real font, after all it's just a number.

The primitive `\setfontid` can be used to enable a font with the given id, which of course needs to be a valid one.

6.4.3 `\noligs` and `\nokerns`

These primitives prohibit ligature and kerning insertion at the time when the initial node list is built by `LUATEX`'s main control loop. You can enable these primitives when you want to do node list processing of 'characters', where `TEX`'s normal processing would get in the way.

```
\noligs <integer>  
\nokerns <integer>
```

These primitives can also be implemented by overloading the ligature building and kerning functions, i.e. by assigning dummy functions to their associated callbacks. Keep in mind that when you define a font (using `LUA`) you can also omit the kern and ligature tables, which has the same effect as the above.

6.4.4 `\nospaces`

This new primitive can be used to overrule the usual `\spaceskip` related heuristics when a space character is seen in a text flow. The value 1 triggers no injection while 2 results in injection of a zero skip. In figure 6.1 we see the results for four characters separated by a space.






Figure 6.1 The `\nospaces` options.

6.4.5 `\protrusionboundary`

The protrusion detection mechanism is enhanced a bit to enable a bit more complex situations. When protrusion characters are identified some nodes are skipped:

- ▶ zero glue
- ▶ penalties
- ▶ empty discretionaries
- ▶ normal zero kerns
- ▶ rules with zero dimensions
- ▶ math nodes with a surround of zero
- ▶ dir nodes
- ▶ empty horizontal lists
- ▶ local par nodes
- ▶ inserts, marks and adjusts
- ▶ boundaries
- ▶ whatsits

Because this can not be enough, you can also use a protrusion boundary node to make the next node being ignored. When the value is 1 or 3, the next node will be ignored in the test when locating a left boundary condition. When the value is 2 or 3, the previous node will be ignored when locating a right boundary condition (the search goes from right to left). This permits protrusion combined with for instance content moved into the margin:

```
\protrusionboundary1\llap{!\quad}«Who needs protrusion?»
```

6.5 The LUA font library

6.5.1 Introduction

The LUA font library is reduced to a few commands. Contrary to $\text{LUA}\text{\TeX}$ there is no loading of TFM or VF files. The explanation of the following commands is in the $\text{LUA}\text{\TeX}$ manual.

FUNCTION	DESCRIPTION
current	returns the id of the currently active font
max	returns the last assigned font identifier
setfont	enables a font setfont (sets the current font id)
addcharacters	adds characters to a font
define	defined a font
id	returns the id that relates to a command name

For practical reasons the management of font identifiers is still done by $\text{T}\text{\TeX}$ but it can become an experiment to delegate that to LUA as well.



6.5.2 Defining a font with `define`, `addcharacters` and `setfont`

Normally you will use a callback to define a font but there's also a LUA function that does the job.

```
id = font.define(<table> f)
```

Within reasonable bounds you can extend a font after it has been defined. Because some properties are best left unchanged this is limited to adding characters.

```
font.addcharacters(<number n>, <table> f)
```

The table passed can have the fields `characters` which is a (sub)table like the one used in `define`, and for virtual fonts a `fonts` table can be added. The characters defined in the `characters` table are added (when not yet present) or replace an existing entry. Keep in mind that replacing can have side effects because a character already can have been used. Instead of posing restrictions we expect the user to be careful. The `setfont` helper is a more drastic replacer and only works when a font has not been used yet.

6.5.3 Font ids: `id`, `max` and `current`

```
<number> i = font.id(<string> csname)
```

This returns the font id associated with `csname`, or `-1` if `csname` is not defined.

```
<number> i = font.max()
```

This is the largest used index so far. The currently active font id can be queried or set with:

```
<number> i = font.current()  
font.current(<number> i)
```





7 Math

7.1 Traditional alongside OPENTYPE

At this point there is no difference between LUAMETATEX and LUATEX with respect to math. The handling of mathematics in LUATEX differs quite a bit from how TEX82 (and therefore PDFTEX) handles math. First, LUATEX adds primitives and extends some others so that UNICODE input can be used easily. Second, all of TEX82 's internal special values (for example for operator spacing) have been made accessible and changeable via control sequences. Third, there are extensions that make it easier to use OPENTYPE math fonts. And finally, there are some extensions that have been proposed or considered in the past that are now added to the engine.

7.2 Unicode math characters

Character handling is now extended up to the full UNICODE range (the $\backslash\text{U}$ prefix), which is compatible with XETEX .

The math primitives from TEX are kept as they are, except for the ones that convert from input to math commands: mathcode , and delcode . These two now allow for a 21-bit character argument on the left hand side of the equals sign.

Some of the new LUATEX primitives read more than one separate value. This is shown in the tables below by a plus sign.

The input for such primitives would look like this:

```
\def\overbrace{\Umathaccent 0 1 "23DE }
```

The altered TEX82 primitives are:

PRIMITIVE	MIN	MAX		MIN	MAX
$\backslash\text{mathcode}$	0	10FFFF	=	0	8000
$\backslash\text{delcode}$	0	10FFFF	=	0	FFFFFF

The unaltered ones are:

PRIMITIVE	MIN	MAX
$\backslash\text{mathchardef}$	0	8000
$\backslash\text{mathchar}$	0	7FFF
$\backslash\text{mathaccent}$	0	7FFF
$\backslash\text{delimiter}$	0	7FFFFFFF
$\backslash\text{radical}$	0	7FFFFFFF

For practical reasons $\backslash\text{mathchardef}$ will silently accept values larger than $0x8000$ and interpret it as $\backslash\text{Umathcharnumdef}$. This is needed to satisfy older macro packages.

The following new primitives are compatible with XETEX :



PRIMITIVE	MIN	MAX	MIN	MAX
\Umathchardef	0+0+0	7+FF+10FFFF		
\Umathcharnumdef ⁵	-80000000	7FFFFFFF		
\Umathcode	0	10FFFF	= 0+0+0	7+FF+10FFFF
\Udelcode	0	10FFFF	= 0+0	FF+10FFFF
\Umathchar	0+0+0	7+FF+10FFFF		
\Umathaccent	0+0+0	7+FF+10FFFF		
\Udelimiter	0+0+0	7+FF+10FFFF		
\Uradical	0+0	FF+10FFFF		
\Umathcharnum	-80000000	7FFFFFFF		
\Umathcodenum	0	10FFFF	= -80000000	7FFFFFFF
\Udelcodenum	0	10FFFF	= -80000000	7FFFFFFF

Specifications typically look like:

```
\Umathchardef\xx="1"0"456
\Umathcode 123="1"0"789
```

The new primitives that deal with delimiter-style objects do not set up a ‘large family’. Selecting a suitable size for display purposes is expected to be dealt with by the font via the `\Umathoperator-size` parameter.

For some of these primitives, all information is packed into a single signed integer. For the first two (`\Umathcharnum` and `\Umathcodenum`), the lowest 21 bits are the character code, the 3 bits above that represent the math class, and the family data is kept in the topmost bits. This means that the values for math families 128–255 are actually negative. For `\Udelcodenum` there is no math class. The math family information is stored in the bits directly on top of the character code. Using these three commands is not as natural as using the two- and three-value commands, so unless you know exactly what you are doing and absolutely require the speedup resulting from the faster input scanning, it is better to use the verbose commands instead.

The `\Umathaccent` command accepts optional keywords to control various details regarding math accents. See section 7.6.2 below for details.

There are more new primitives and all of these will be explained in following sections:

PRIMITIVE	VALUE RANGE (IN HEX)
\Uroot	0 + 0-FF + 10FFFF
\Uoverdelimiter	0 + 0-FF + 10FFFF
\Uunderdelimiter	0 + 0-FF + 10FFFF
\Udelimiterover	0 + 0-FF + 10FFFF
\Udelimiterunder	0 + 0-FF + 10FFFF

7.3 Math styles

7.3.1 \mathstyle

It is possible to discover the math style that will be used for a formula in an expandable fashion (while the math list is still being read). To make this possible, L^AT_EX adds the new primitive:



`\mathstyle`. This is a ‘convert command’ like e.g. `\romannumeral`: its value can only be read, not set.

The returned value is between 0 and 7 (in math mode), or -1 (all other modes). For easy testing, the eight math style commands have been altered so that they can be used as numeric values, so you can write code like this:

```
\ifnum\mathstyle=\textstyle
  \message{normal text style}
\else \ifnum\mathstyle=\crampedtextstyle
  \message{cramped text style}
\fi \fi
```

Sometimes you won’t get what you expect so a bit of explanation might help to understand what happens. When math is parsed and expanded it gets turned into a linked list. In a second pass the formula will be build. This has to do with the fact that in order to determine the automatically chosen sizes (in for instance fractions) following content can influence preceding sizes. A side effect of this is for instance that one cannot change the definition of a font family (and thereby reusing numbers) because the number that got used is stored and used in the second pass (so changing `\fam 12` mid-formula spoils over to preceding use of that family).

The style switching primitives like `\textstyle` are turned into nodes so the styles set there are frozen. The `\mathchoice` primitive results in four lists being constructed of which one is used in the second pass. The fact that some automatic styles are not yet known also means that the `\mathstyle` primitive expands to the current style which can of course be different from the one really used. It’s a snapshot of the first pass state. As a consequence in the following example you get a style number (first pass) typeset that can actually differ from the used style (second pass). In the case of a math choice used ungrouped, the chosen style is used after the choice too, unless you group.

```
[a:\mathstyle]\quad
\bgroup
\mathchoice
  {\bf \scriptstyle      (x:d :\mathstyle)}
  {\bf \scriptscriptstyle (x:t :\mathstyle)}
  {\bf \scriptscriptstyle (x:s :\mathstyle)}
  {\bf \scriptscriptstyle (x:ss:\mathstyle)}
\egroup
\quad[b:\mathstyle]\quad
\mathchoice
  {\bf \scriptstyle      (y:d :\mathstyle)}
  {\bf \scriptscriptstyle (y:t :\mathstyle)}
  {\bf \scriptscriptstyle (y:s :\mathstyle)}
  {\bf \scriptscriptstyle (y:ss:\mathstyle)}
\quad[c:\mathstyle]\quad
\bgroup
\mathchoice
  {\bf \scriptstyle      (z:d :\mathstyle)}
  {\bf \scriptscriptstyle (z:t :\mathstyle)}
```



```

{\bf \scriptscriptstyle (z:s:\mathstyle)}
{\bf \scriptscriptstyle (z:ss:\mathstyle)}
\egroup
\quad[d:\mathstyle]

```

This gives:

```
[a:0] (x:d:4) [b:0] (y:d:4) [c:0] (z:s:6) [d:0]
```

```
[a:2] (x:t:6) [b:2] (y:t:6) [c:2] (z:ss:6) [d:2]
```

Using `\begingroup ... \endgroup` instead gives:

```
[a:0] (x:d:4) [b:0] (y:s:6) [c:0] (z:ss:6) [d:0]
```

```
[a:2] (x:t:6) [b:2] (y:ss:6) [c:2] (z:ss:6) [d:2]
```

This might look wrong but it's just a side effect of `\mathstyle` expanding to the current (first pass) style and the number being injected in the list that gets converted in the second pass. It all makes sense and it illustrates the importance of grouping. In fact, the math choice style being effective afterwards has advantages. It would be hard to get it otherwise.

7.3.2 `\Ustack`

There are a few math commands in \TeX where the style that will be used is not known straight from the start. These commands (`\over`, `\atop`, `\overwithdelims`, `\atopwithdelims`) would therefore normally return wrong values for `\mathstyle`. To fix this, $\text{LUA}\TeX$ introduces a special prefix command: `\Ustack`:

```
\Ustack {a \over b}$
```

The `\Ustack` command will scan the next brace and start a new math group with the correct (numerator) math style.

7.3.3 The new `\cramped ... style` commands

$\text{LUA}\TeX$ has four new primitives to set the cramped math styles directly:

```

\crampeddisplaystyle
\crampedtextstyle
\crampedscriptstyle
\crampedscriptscriptstyle

```

These additional commands are not all that valuable on their own, but they come in handy as arguments to the math parameter settings that will be added shortly.

In Eijkhouts “ \TeX by Topic” the rules for handling styles in scripts are described as follows:

- ▶ In any style superscripts and subscripts are taken from the next smaller style. Exception: in display style they are in script style.
- ▶ Subscripts are always in the cramped variant of the style; superscripts are only cramped if the original style was cramped.



- ▶ In an `.. \over ..` formula in any style the numerator and denominator are taken from the next smaller style.
- ▶ The denominator is always in cramped style; the numerator is only in cramped style if the original style was cramped.
- ▶ Formulas under a `\sqrt` or `\overline` are in cramped style.

In L^AT_EX one can set the styles in more detail which means that you sometimes have to set both normal and cramped styles to get the effect you want. (Even) if we force styles in the script using `\scriptstyle` and `\crampedscriptstyle` we get this:

STYLE	EXAMPLE
default	$b^{x=xx}$ $x=xx$
script	$b^{x=xx}$ $x=xx$
crampedscript	$b^{x=xx}$ $x=xx$

Now we set the following parameters

```
\Umathordrelspacing\scriptstyle=30mu
\Umathordordspacing\scriptstyle=30mu
```

This gives a different result:

STYLE	EXAMPLE
default	$b^x = x \quad x$ $x=xx$
script	$b^x = x \quad x$ $x = x \quad x$
crampedscript	$b^{x=xx}$ $x=xx$

But, as this is not what is expected (visually) we should say:

```
\Umathordrelspacing\scriptstyle=30mu
\Umathordordspacing\scriptstyle=30mu
\Umathordrelspacing\crampedscriptstyle=30mu
\Umathordordspacing\crampedscriptstyle=30mu
```

Now we get:

STYLE	EXAMPLE
default	$b^x = x \quad x$ $x = x \quad x$
script	$b^x = x \quad x$ $x = x \quad x$
crampedscript	$b^x = x \quad x$ $x = x \quad x$

7.4 Math parameter settings

7.4.1 Many new `\Umath*` primitives

In L^AT_EX, the font dimension parameters that T_EX used in math typesetting are now accessible via primitive commands. In fact, refactoring of the math engine has resulted in many more parameters than were not accessible before.



PRIMITIVE NAME	DESCRIPTION
<code>\Umathquad</code>	the width of 18 mu's
<code>\Umathaxis</code>	height of the vertical center axis of the math formula above the baseline
<code>\Umathoperatorsize</code>	minimum size of large operators in display mode
<code>\Umathoverbarkern</code>	vertical clearance above the rule
<code>\Umathoverbarrule</code>	the width of the rule
<code>\Umathoverbarvgap</code>	vertical clearance below the rule
<code>\Umathunderbarkern</code>	vertical clearance below the rule
<code>\Umathunderbarrule</code>	the width of the rule
<code>\Umathunderbarvgap</code>	vertical clearance above the rule
<code>\Umathradicalkern</code>	vertical clearance above the rule
<code>\Umathradicalrule</code>	the width of the rule
<code>\Umathradicalvgap</code>	vertical clearance below the rule
<code>\Umathradicaldegreebefore</code>	the forward kern that takes place before placement of the radical degree
<code>\Umathradicaldegreeafter</code>	the backward kern that takes place after placement of the radical degree
<code>\Umathradicaldegreeraise</code>	this is the percentage of the total height and depth of the radical sign that the degree is raised by; it is expressed in percents, so 60% is expressed as the integer 60
<code>\Umathstackvgap</code>	vertical clearance between the two elements in a <code>\atop</code> stack
<code>\Umathstacknumup</code>	numerator shift upward in <code>\atop</code> stack
<code>\Umathstackdenomdown</code>	denominator shift downward in <code>\atop</code> stack
<code>\Umathfractionrule</code>	the width of the rule in a <code>\over</code>
<code>\Umathfractionnumvgap</code>	vertical clearance between the numerator and the rule
<code>\Umathfractionnumup</code>	numerator shift upward in <code>\over</code>
<code>\Umathfractiondenomvgap</code>	vertical clearance between the denominator and the rule
<code>\Umathfractiondenomdown</code>	denominator shift downward in <code>\over</code>
<code>\Umathfractiondelsize</code>	minimum delimiter size for <code>\dots</code> withdelims
<code>\Umathlimitabovevgap</code>	vertical clearance for limits above operators
<code>\Umathlimitabovebgap</code>	vertical baseline clearance for limits above operators
<code>\Umathlimitabovekern</code>	space reserved at the top of the limit
<code>\Umathlimitbelowvgap</code>	vertical clearance for limits below operators
<code>\Umathlimitbelowbgap</code>	vertical baseline clearance for limits below operators
<code>\Umathlimitbelowkern</code>	space reserved at the bottom of the limit
<code>\Umathoverdelimitervgap</code>	vertical clearance for limits above delimiters
<code>\Umathoverdelimiterbgap</code>	vertical baseline clearance for limits above delimiters
<code>\Umathunderdelimitervgap</code>	vertical clearance for limits below delimiters
<code>\Umathunderdelimiterbgap</code>	vertical baseline clearance for limits below delimiters
<code>\Umathsubshiftdrop</code>	subscript drop for boxes and subformulas
<code>\Umathsubshiftdown</code>	subscript drop for characters
<code>\Umathsupshiftdrop</code>	superscript drop (raise, actually) for boxes and subformulas
<code>\Umathsupshiftpup</code>	superscript raise for characters
<code>\Umathsubsupshiftdown</code>	subscript drop in the presence of a superscript



<code>\Umathsubtopmax</code>	the top of standalone subscripts cannot be higher than this above the baseline
<code>\Umathsupbottommin</code>	the bottom of standalone superscripts cannot be less than this above the baseline
<code>\Umathsupsubbottommax</code>	the bottom of the superscript of a combined super- and subscript be at least as high as this above the baseline
<code>\Umathsubsupvgap</code>	vertical clearance between super- and subscript
<code>\Umathspaceafterscript</code>	additional space added after a super- or subscript
<code>\Umathconnectoroverlapmin</code>	minimum overlap between parts in an extensible recipe

Each of the parameters in this section can be set by a command like this:

```
\Umathquad\displaystyle=1em
```

they obey grouping, and you can use `\the\Umathquad\displaystyle` if needed.

7.4.2 Font-based math parameters

While it is nice to have these math parameters available for tweaking, it would be tedious to have to set each of them by hand. For this reason, L^AT_EX initializes a bunch of these parameters whenever you assign a font identifier to a math family based on either the traditional math font dimensions in the font (for assignments to math family 2 and 3 using TFM-based fonts like `cmsy` and `cmex`), or based on the named values in a potential `MathConstants` table when the font is loaded via Lua. If there is a `MathConstants` table, this takes precedence over font dimensions, and in that case no attention is paid to which family is being assigned to: the `MathConstants` tables in the last assigned family sets all parameters.

In the table below, the one-letter style abbreviations and symbolic tfm font dimension names match those used in the T_EXbook. Assignments to `\textfont` set the values for the cramped and uncramped display and text styles, `\scriptfont` sets the script styles, and `\scriptscriptfont` sets the scriptscript styles, so we have eight parameters for three font sizes. In the TFM case, assignments only happen in family 2 and family 3 (and of course only for the parameters for which there are font dimensions).

Besides the parameters below, L^AT_EX also looks at the ‘space’ font dimension parameter. For math fonts, this should be set to zero.

VARIABLE / STYLE	TFM / OPENTYPE
<code>\Umathaxis</code>	<code>axis_height</code> <code>AxisHeight</code>
⁶ <code>\Umathoperatorsiz</code> D, D'	— <code>DisplayOperatorMinHeight</code>
⁹ <code>\Umathfractiondelsize</code> D, D'	<code>delim1</code> <code>FractionDelimiterDisplayStyleSize</code>
⁹ <code>\Umathfractiondelsize</code> T, T', S, S', SS, SS'	<code>delim2</code> <code>FractionDelimiterSize</code>
<code>\Umathfractiondenomdown</code>	<code>denom1</code>



D, D'	FractionDenominatorDisplayStyleShiftDown
<code>\Umathfractiondenomdown</code>	denom2
T, T', S, S', SS, SS'	FractionDenominatorShiftDown
<code>\Umathfractiondenomvgap</code>	3*default_rule_thickness
D, D'	FractionDenominatorDisplayStyleGapMin
<code>\Umathfractiondenomvgap</code>	default_rule_thickness
T, T', S, S', SS, SS'	FractionDenominatorGapMin
<code>\Umathfractionnumup</code>	num1
D, D'	FractionNumeratorDisplayStyleShiftUp
<code>\Umathfractionnumup</code>	num2
T, T', S, S', SS, SS'	FractionNumeratorShiftUp
<code>\Umathfractionnumvgap</code>	3*default_rule_thickness
D, D'	FractionNumeratorDisplayStyleGapMin
<code>\Umathfractionnumvgap</code>	default_rule_thickness
T, T', S, S', SS, SS'	FractionNumeratorGapMin
<code>\Umathfractionrule</code>	default_rule_thickness
	FractionRuleThickness
<code>\Umathskewedfractionhgap</code>	math_quad/2
	SkewedFractionHorizontalGap
<code>\Umathskewedfractionvgap</code>	math_x_height
	SkewedFractionVerticalGap
<code>\Umathlimitabovebgap</code>	big_op_spacing3
	UpperLimitBaselineRiseMin
¹ <code>\Umathlimitabovekern</code>	big_op_spacing5
	0
<code>\Umathlimitabovevgap</code>	big_op_spacing1
	UpperLimitGapMin
<code>\Umathlimitbelowbgap</code>	big_op_spacing4
	LowerLimitBaselineDropMin
¹ <code>\Umathlimitbelowkern</code>	big_op_spacing5
	0
<code>\Umathlimitbelowvgap</code>	big_op_spacing2
	LowerLimitGapMin
<code>\Umathoverdelimitervgap</code>	big_op_spacing1
	StretchStackGapBelowMin
<code>\Umathoverdelimiterbgap</code>	big_op_spacing3
	StretchStackTopShiftUp
<code>\Umathunderdelimitervgap</code>	big_op_spacing2
	StretchStackGapAboveMin
<code>\Umathunderdelimiterbgap</code>	big_op_spacing4
	StretchStackBottomShiftDown
<code>\Umathoverbarkern</code>	default_rule_thickness



	OverbarExtraAscender
\Umathoverbarrule	default_rule_thickness OverbarRuleThickness
\Umathoverbarvgap	3*default_rule_thickness OverbarVerticalGap
¹ \Umathquad	math_quad <font_size(f)>
\Umathradicalkern	default_rule_thickness RadicalExtraAscender
² \Umathradicalrule	<not set> RadicalRuleThickness
³ \Umathradicalvgap D, D'	default_rule_thickness+abs(math_x_height)/4 RadicalDisplayStyleVerticalGap
³ \Umathradicalvgap T, T', S, S', SS, SS'	default_rule_thickness+abs(default_rule_thickness)/4 RadicalVerticalGap
² \Umathradicaldegreebefore	<not set> RadicalKernBeforeDegree
² \Umathradicaldegreeafter	<not set> RadicalKernAfterDegree
^{2,7} \Umathradicaldegreeraise	<not set> RadicalDegreeBottomRaisePercent
⁴ \Umathspaceafterscript	script_space SpaceAfterScript
\Umathstackdenomdown D, D'	denom1 StackBottomDisplayStyleShiftDown
\Umathstackdenomdown T, T', S, S', SS, SS'	denom2 StackBottomShiftDown
\Umathstacknumup D, D'	num1 StackTopDisplayStyleShiftUp
\Umathstacknumup T, T', S, S', SS, SS'	num3 StackTopShiftUp
\Umathstackvgap D, D'	7*default_rule_thickness StackDisplayStyleGapMin
\Umathstackvgap T, T', S, S', SS, SS'	3*default_rule_thickness StackGapMin
\Umathsubshiftdown	sub1 SubscriptShiftDown
\Umathsubshiftdrop	sub_drop SubscriptBaselineDropMin
⁸ \Umathsubsupshiftdown	— SubscriptShiftDownWithSuperscript
\Umathsubtopmax	abs(math_x_height*4)/5



	SubscriptTopMax
\Umathsubsupvgap	4*default_rule_thickness SubSuperscriptGapMin
\Umathsupbottommin	abs(math_x_height/4) SuperscriptBottomMin
\Umathsupshiftdrop	sup_drop SuperscriptBaselineDropMax
\Umathsupshiftup D	sup1 SuperscriptShiftUp
\Umathsupshiftup T, S, SS,	sup2 SuperscriptShiftUp
\Umathsupshiftup D', T', S', SS'	sup3 SuperscriptShiftUpCramped
\Umathsupsubbottommax	abs(math_x_height*4)/5 SuperscriptBottomMaxWithSubscript
\Umathunderbarkern	default_rule_thickness UnderbarExtraDescender
\Umathunderbarrule	default_rule_thickness UnderbarRuleThickness
\Umathunderbarvgap	3*default_rule_thickness UnderbarVerticalGap
⁵ \Umathconnectoroverlapmin	0 MinConnectorOverlap

Note 1: OPENTYPE fonts set `\Umathlimitabovekern` and `\Umathlimitbelowkern` to zero and set `\Umathquad` to the font size of the used font, because these are not supported in the MATH table,

Note 2: Traditional TFM fonts do not set `\Umathradicalrule` because $\mathrm{T}_{\mathrm{E}}\mathrm{X}82$ uses the height of the radical instead. When this parameter is indeed not set when $\mathrm{L}^{\mathrm{A}}\mathrm{T}_{\mathrm{E}}\mathrm{X}$ has to typeset a radical, a backward compatibility mode will kick in that assumes that an oldstyle $\mathrm{T}_{\mathrm{E}}\mathrm{X}$ font is used. Also, they do not set `\Umathradicaldegreebefore`, `\Umathradicaldegreeafter`, and `\Umathradicaldegreerise`. These are then automatically initialized to $5/18\mathrm{quad}$, $-10/18\mathrm{quad}$, and 60.

Note 3: If TFM fonts are used, then the `\Umathradicalvgap` is not set until the first time $\mathrm{L}^{\mathrm{A}}\mathrm{T}_{\mathrm{E}}\mathrm{X}$ has to typeset a formula because this needs parameters from both family 2 and family 3. This provides a partial backward compatibility with $\mathrm{T}_{\mathrm{E}}\mathrm{X}82$, but that compatibility is only partial: once the `\Umathradicalvgap` is set, it will not be recalculated any more.

Note 4: When TFM fonts are used a similar situation arises with respect to `\Umathspaceafterscript`: it is not set until the first time $\mathrm{L}^{\mathrm{A}}\mathrm{T}_{\mathrm{E}}\mathrm{X}$ has to typeset a formula. This provides some backward compatibility with $\mathrm{T}_{\mathrm{E}}\mathrm{X}82$. But once the `\Umathspaceafterscript` is set, `\scriptspace` will never be looked at again.

Note 5: Traditional TFM fonts set `\Umathconnectoroverlapmin` to zero because $\mathrm{T}_{\mathrm{E}}\mathrm{X}82$ always stacks extensibles without any overlap.



Note 6: The `\Umathoperator size` is only used in `\displaystyle`, and is only set in OPENTYPE fonts. In TFM font mode, it is artificially set to one scaled point more than the initial attempt's size, so that always the 'first next' will be tried, just like in T_EX82.

Note 7: The `\Umathradicaldegreeraise` is a special case because it is the only parameter that is expressed in a percentage instead of a number of scaled points.

Note 8: `SubscriptShiftDownWithSuperscript` does not actually exist in the 'standard' OPENTYPE math font Cambria, but it is useful enough to be added.

Note 9: `FractionDelimiterDisplayStyleSize` and `FractionDelimiterSize` do not actually exist in the 'standard' OPENTYPE math font Cambria, but were useful enough to be added.

7.5 Math spacing

7.5.1 Setting inline surrounding space with `\mathsurround[skip]\mathsurround[skip]`

Inline math is surrounded by (optional) `\mathsurround` spacing but that is a fixed dimension. There is now an additional parameter `\mathsurroundskip`. When set to a non-zero value (or zero with some stretch or shrink) this parameter will replace `\mathsurround`. By using an additional parameter instead of changing the nature of `\mathsurround`, we can remain compatible. In the meantime a bit more control has been added via `\mathsurroundmode`. This directive can take 6 values with zero being the default behaviour.

```
\mathsurround      10pt
\mathsurroundskip20pt
```

MODE	X\$X\$X	X \$X\$ X	EFFECT
0	xxx	x x x	obey <code>\mathsurround</code> when <code>\mathsurroundskip</code> is 0pt
1	xxx	x x x	only add skip to the left
2	xxx	x x x	only add skip to the right
3	xxx	x x x	add skip to the left and right
4	xxx	x x x	ignore the skip setting, obey <code>\mathsurround</code>
5	xxx	x x x	disable all spacing around math
6	xxx	x x x	only apply <code>\mathsurroundskip</code> when also spacing
7	xxx	x x x	only apply <code>\mathsurroundskip</code> when no spacing

Method six omits the surround glue when there is (x)spacing glue present while method seven does the opposite, the glue is only applied when there is (x)space glue present too. Anything more fancy, like checking the beginning or end of a paragraph (or edges of a box) would not be robust anyway. If you want that you can write a callback that runs over a list and analyzes a paragraph. Actually, in that case you could also inject glue (or set the properties of a math node) explicitly. So, these modes are in practice mostly useful for special purposes and experiments (they originate in a tracker item). Keep in mind that this glue is part of the math node and not always treated as normal glue: it travels with the begin and end math nodes. Also, method 6 and 7 will zero the skip related fields in a node when applicable in the first occasion that checks them (linebreaking or packaging).



7.5.2 Pairwise spacing and `\Umath...spacing` commands

Besides the parameters mentioned in the previous sections, there are also 64 new primitives to control the math spacing table (as explained in Chapter 18 of the `TEXbook`). The primitive names are a simple matter of combining two math atom types, but for completeness' sake, here is the whole list:

<code>\Umathordordspacing</code>	<code>\Umathopenordspacing</code>
<code>\Umathordopspacing</code>	<code>\Umathopenopspacing</code>
<code>\Umathordbinspacing</code>	<code>\Umathopenbinspacing</code>
<code>\Umathordrelspacing</code>	<code>\Umathopenrelspacing</code>
<code>\Umathordopenspacing</code>	<code>\Umathopenopenspacing</code>
<code>\Umathordclosespacing</code>	<code>\Umathopenclosespacing</code>
<code>\Umathordpunctspacing</code>	<code>\Umathopenpunctspacing</code>
<code>\Umathordinnerspacing</code>	<code>\Umathopeninnerspacing</code>
<code>\Umathopordspacing</code>	<code>\Umathcloseordspacing</code>
<code>\Umathopopspacing</code>	<code>\Umathcloseopspacing</code>
<code>\Umathopbinspacing</code>	<code>\Umathclosebinspacing</code>
<code>\Umathoprelspacing</code>	<code>\Umathclosere-spacing</code>
<code>\Umathopopenspacing</code>	<code>\Umathcloseopenspacing</code>
<code>\Umathopclosespacing</code>	<code>\Umathcloseclosespacing</code>
<code>\Umathoppunctspacing</code>	<code>\Umathclosepunctspacing</code>
<code>\Umathopinnerspacing</code>	<code>\Umathcloseinnerspacing</code>
<code>\Umathbinordspacing</code>	<code>\Umathpunctordspacing</code>
<code>\Umathbinopspacing</code>	<code>\Umathpunctopspacing</code>
<code>\Umathbinbinspacing</code>	<code>\Umathpunctbinspacing</code>
<code>\Umathbinrelspacing</code>	<code>\Umathpunctrelspacing</code>
<code>\Umathbinopenspacing</code>	<code>\Umathpunctopenspacing</code>
<code>\Umathbinclosespacing</code>	<code>\Umathpunctclosespacing</code>
<code>\Umathbinpunctspacing</code>	<code>\Umathpunctpunctspacing</code>
<code>\Umathbininnerspacing</code>	<code>\Umathpunctinnerspacing</code>
<code>\Umathrelordspacing</code>	<code>\Umathinnerordspacing</code>
<code>\Umathrelopspacing</code>	<code>\Umathinneropspacing</code>
<code>\Umathrelbinspacing</code>	<code>\Umathinnerbinspacing</code>
<code>\Umathrelrelspacing</code>	<code>\Umathinnerrelspacing</code>
<code>\Umathrelopenspacing</code>	<code>\Umathinneropenspacing</code>
<code>\Umathrelclosespacing</code>	<code>\Umathinnerclosespacing</code>
<code>\Umathrelpunctspacing</code>	<code>\Umathinnerpunctspacing</code>
<code>\Umathrelinnerspacing</code>	<code>\Umathinnerinnerspacing</code>

These parameters are of type `\muskip`, so setting a parameter can be done like this:

```
\Umathopordspacing\displaystyle=4mu plus 2mu
```

They are all initialized by `initex` to the values mentioned in the table in Chapter 18 of the `TEXbook`.



Note 1: for ease of use as well as for backward compatibility, `\thinmuskip`, `\medmuskip` and `\thickmuskip` are treated specially. In their case a pointer to the corresponding internal parameter is saved, not the actual `\muskip` value. This means that any later changes to one of these three parameters will be taken into account.

Note 2: Careful readers will realise that there are also primitives for the items marked * in the `TEXbook`. These will not actually be used as those combinations of atoms cannot actually happen, but it seemed better not to break orthogonality. They are initialized to zero.

7.5.3 Local settings

Math is processed in two passes. The first pass is needed to intercept for instance `\over`, one of the few `TEX` commands that actually has a preceding argument. There are often lots of curly braces used in math and these can result in a nested run of the math sub engine. However, you need to be aware of the fact that some properties are kind of global to a formula and the last setting (for instance a family switch) wins. This also means that a change (or again, the last one) in math parameters affects the whole formula. In `LUAMETATEX` we have changed this model a bit. One can argue that this introduces an incompatibility but it's hard to imagine a reason for setting the parameters at the end of a formula run and assume that they also influence what goes in front.

```
$
                                x \Usubscript {-}
\frozen\Umathsubshiftdown\textstyle 0pt x \Usubscript {0}
{\frozen\Umathsubshiftdown\textstyle 5pt x \Usubscript {5}}
                                x \Usubscript {0}
{\frozen\Umathsubshiftdown\textstyle 15pt x \Usubscript {15}}
                                x \Usubscript {0}
{\frozen\Umathsubshiftdown\textstyle 20pt x \Usubscript {20}}
                                x \Usubscript {0}
\frozen\Umathsubshiftdown\textstyle 10pt x \Usubscript {10}
                                x \Usubscript {0}
$
```

The `\frozen` prefix does the magic: it injects information in the math list about the set parameter. In `LUATEX 1.10` the last setting, the 10pt drop wins, but in `LUAMETATEX` you will see each local setting taking effect. The implementation uses a new node type, parameters nodes, so you you might encounter these in an unprocessed math list. The result looks as follows:

```
x_x0x_5x0x  x0x  x0x  x
      15      20      10 0
```

7.5.4 Skips around display math and `\mathdisplayskipmode`

The injection of `\abovedisplayskip` and `\belowdisplayskip` is not symmetrical. An above one is always inserted, also when zero, but the below is only inserted when larger than zero. Especially the latter makes it sometimes hard to fully control spacing. Therefore `LUATEX` comes with a new directive: `\mathdisplayskipmode`. The following values apply:



VALUE	MEANING
0	normal behaviour
1	always (same as 0)
2	only when not zero
3	never, not even when not zero

7.5.5 Nolimit correction with `\mathnolimitsmode`

There are two extra math parameters `\Umathnolimitsupfactor` and `\Umathnolimitssubfactor` that were added to provide some control over how limits are spaced (for example the position of super and subscripts after integral operators). They relate to an extra parameter `\mathnolimitsmode`. The half corrections are what happens when scripts are placed above and below. The problem with italic corrections is that officially that correction italic is used for above/below placement while advanced kerns are used for placement at the right end. The question is: how often is this implemented, and if so, do the kerns assume correction too. Anyway, with this parameter one can control it.

	\int_1^0	\int_1^0	\int_1^0	\int_1^0	\int_1^0	\int_1^0
mode	0	1	2	3	4	8000
superscript	0	font	0	0	+ic/2	0
subscript	-ic	font	0	-ic/2	-ic/2	8000ic/1000

When the mode is set to one, the math parameters are used. This way a macro package writer can decide what looks best. Given the current state of fonts in `CONTEX`T we currently use mode 1 with factor 0 for the superscript and 750 for the subscripts. Positive values are used for both parameters but the subscript shifts to the left. A `\mathnolimitsmode` larger than 15 is considered to be a factor for the subscript correction. This feature can be handy when experimenting.

7.5.6 Controlling math italic mess with `\mathitalicsmode`

The `\mathitalicsmode` parameter can be set to 1 to force italic correction before noads that represent some more complex structure (read: everything that is not an ord, bin, rel, open, close, punct or inner). We show a Cambria example.

```
\mathitalicsmode = 0  [T1] [T] [T+1] [T21] [T√1]
\mathitalicsmode = 1  [T1] [T] [T+1] [T21] [T√1]
```

This kind of parameters relate to the fact that italic correction in `OPENTYPE` math is bound to fuzzy rules. So, control is the solution.

7.5.7 Influencing script kerning with `\mathscriptboxmode`

If you want to typeset text in math macro packages often provide something `\text` which obeys the script sizes. As the definition can be anything there is a good chance that the kerning doesn't come out well when used in a script. Given that the first glyph ends up in a `\hbox` we have some



control over this. And, as a bonus we also added control over the normal sublist kerning. The `\mathscriptboxmode` parameter defaults to 1.

VALUE	MEANING
0	forget about kerning
1	kern math sub lists with a valid glyph
2	also kern math sub boxes that have a valid glyph
2	only kern math sub boxes with a boundary node present

Here we show some examples. Of course this doesn't solve all our problems, if only because some fonts have characters with bounding boxes that compensate for italics, while other fonts can lack kerns.

	<code>\$T_{\text{tf fluff}}\$</code>	<code>\$T_{\text{tf fluff}}\$</code>	<code>\$T_{\text{\text{fluff}}}\$</code>	<code>\$T_{\text{\text{fluff}}}\$</code>	<code>\$T_{\text{\text{boundary1 fluff}}}\$</code>
	mode 0	mode 1	mode 1	mode 2	mode 3
modern	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}
lucidaot	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}
pagella	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}
cambria	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}
dejavu	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}	T_{fluff}

Kerning between a character subscript is controlled by `\mathscriptcharmode` which also defaults to 1.

Here is another example. Internally we tag kerns as italic kerns or font kerns where font kerns result from the staircase kern tables. In 2018 fonts like Latin Modern and Pagella rely on cheats with the boundingbox, Cambria uses staircase kerns and Lucida a mixture. Depending on how fonts evolve we might add some more control over what one can turn on and off.

normal	modern	T_f	γ_e	γ_{ee}	T_{fluff}
	pagella	T_f	γ_e	γ_{ee}	T_{fluff}
	cambria	T_f	γ_e	γ_{ee}	T_{fluff}
	lucidaot	T_f	γ_e	γ_{ee}	T_{fluff}
bold	modern	T_f	γ_e	γ_{ee}	T_{fluff}
	pagella	T_f	γ_e	γ_{ee}	T_{fluff}
	cambria	T_f	γ_e	γ_{ee}	T_{fluff}
	lucidaot	T_f	γ_e	γ_{ee}	T_{fluff}

7.5.8 Forcing fixed scripts with `\mathscriptsmode`

We have three parameters that are used for this fixed anchoring:

PARAMETER	REGISTER
<code>d</code>	<code>\Umathsubshiftdown</code>



u	<code>\Umathsupshiftup</code>
s	<code>\Umathsubsupshiftdown</code>

When we set `\mathscriptsmode` to a value other than zero these are used for calculating fixed positions. This is something that is needed for instance for chemistry. You can manipulate the mentioned variables to achieve different effects.

MODE	DOWN	UP	EXAMPLE
0	dynamic	dynamic	$\text{CH}_2 + \text{CH}_2^+ + \text{CH}_2^2$
1	d	u	$\text{CH}_2 + \text{CH}_2^+ + \text{CH}_2^2$
2	s	u	$\text{CH}_2 + \text{CH}_2^+ + \text{CH}_2^2$
3	s	$u + s - d$	$\text{CH}_2 + \text{CH}_2^+ + \text{CH}_2^2$
4	$d + (s - d)/2$	$u + (s - d)/2$	$\text{CH}_2 + \text{CH}_2^+ + \text{CH}_2^2$
5	d	$u + s - d$	$\text{CH}_2 + \text{CH}_2^+ + \text{CH}_2^2$

The value of this parameter obeys grouping but applies to the whole current formula.

7.5.9 Penalties: `\mathpenaltiesmode`

Only in inline math penalties will be added in a math list. You can force penalties (also in display math) by setting:

```
\mathpenaltiesmode = 1
```

This primitive is not really needed in L^AT_EX because you can use the callback `mlist_to_hlist` to force penalties by just calling the regular routine with forced penalties. However, as part of opening up and control this primitive makes sense. As a bonus we also provide two extra penalties:

```
\prebinoppenalty = -100 % example value
\prerelpenalty   =  900 % example value
```

They default to infinite which signals that they don't need to be inserted. When set they are injected before a binop or rel node. This is an experimental feature.

7.5.10 Equation spacing: `\matheqnogapstep`

By default T_EX will add one quad between the equation and the number. This is hard coded. A new primitive can control this:

```
\matheqnogapstep = 1000
```

Because a math quad from the math text font is used instead of a dimension, we use a step to control the size. A value of zero will suppress the gap. The step is divided by 1000 which is the usual way to mimick floating point factors in T_EX.



7.6 Math constructs

7.6.1 Unscaled fences and `\mathdelimitersmode`

The `\mathdelimitersmode` primitive is experimental and deals with the following (potential) problems. Three bits can be set. The first bit prevents an unwanted shift when the fence symbol is not scaled (a Cambria side effect). The second bit forces italic correction between a preceding character ordinal and the fenced subformula, while the third bit turns that subformula into an ordinary so that the same spacing applies as with unfenced variants. Here we show Cambria (with `\mathitalicsmode` enabled).

<code>\mathdelimitersmode = 0</code>		
<code>\mathdelimitersmode = 1</code>		
<code>\mathdelimitersmode = 2</code>		
<code>\mathdelimitersmode = 3</code>		
<code>\mathdelimitersmode = 4</code>		
<code>\mathdelimitersmode = 5</code>		
<code>\mathdelimitersmode = 6</code>		
<code>\mathdelimitersmode = 7</code>		

So, when set to 7 fenced subformulas with unscaled delimiters come out the same as unfenced ones. This can be handy for cases where one is forced to use `\left` and `\right` always because of unpredictable content. As said, it's an experimental feature (which somehow fits in the exceptional way fences are dealt with in the engine). The full list of flags is given in the next table:

VALUE	MEANING
"01	don't apply the usual shift
"02	apply italic correction when possible
"04	force an ordinary subformula
"08	no shift when a base character
"10	only shift when an extensible

The effect can depend on the font (and for Cambria one can use for instance "16).

7.6.2 Accent handling with `\Umathaccent`

L^AT_EX supports both top accents and bottom accents in math mode, and math accents stretch automatically (if this is supported by the font the accent comes from, of course). Bottom and combined accents as well as fixed-width math accents are controlled by optional keywords following `\Umathaccent`.

The keyword `bottom` after `\Umathaccent` signals that a bottom accent is needed, and the keyword `both` signals that both a top and a bottom accent are needed (in this case two accents need to be specified, of course).

Then the set of three integers defining the accent is read. This set of integers can be prefixed by the fixed keyword to indicate that a non-stretching variant is requested (in case of both accents, this step is repeated).



A simple example:

```
\Umathaccent both fixed 0 0 "20D7 fixed 0 0 "20D7 {example}
```

If a math top accent has to be placed and the accentee is a character and has a non-zero `top_accent` value, then this value will be used to place the accent instead of the `\skewchar` kern used by \TeX 82.

The `top_accent` value represents a vertical line somewhere in the accentee. The accent will be shifted horizontally such that its own `top_accent` line coincides with the one from the accentee. If the `top_accent` value of the accent is zero, then half the width of the accent followed by its italic correction is used instead.

The vertical placement of a top accent depends on the `x_height` of the font of the accentee (as explained in the \TeX book), but if a value turns out to be zero and the font had a `MathConstants` table, then `AccentBaseHeight` is used instead.

The vertical placement of a bottom accent is straight below the accentee, no correction takes place.

Possible locations are `top`, `bottom`, `both` and `center`. When no location is given `top` is assumed. An additional parameter `fraction` can be specified followed by a number; a value of for instance 1200 means that the criterium is 1.2 times the width of the nucleus. The `fraction` only applies to the stepwise selected shapes and is mostly meant for the `overlay` location. It also works for the other locations but then it concerns the width.

7.6.3 Building radicals with `\Uradical` and `\Uroot`

The new primitive `\Uroot` allows the construction of a radical noad including a degree field. Its syntax is an extension of `\Uradical`:

```
\Uradical <fam integer> <char integer> <radicand>  
\Uroot    <fam integer> <char integer> <degree> <radicand>
```

The placement of the degree is controlled by the math parameters `\Umathradicaldegreebefore`, `\Umathradicaldegreeafter`, and `\Umathradicaldegreeraise`. The degree will be typeset in `\scriptscriptstyle`.

7.6.4 Super- and subscripts

The character fields in a LUA-loaded OPENTYPE math font can have a ‘`mathkern`’ table. The format of this table is the same as the ‘`mathkern`’ table that is returned by the `fontloader` library, except that all height and kern values have to be specified in actual scaled points.

When a super- or subscript has to be placed next to a math item, $\text{LUA}\TeX$ checks whether the super- or subscript and the nucleus are both simple character items. If they are, and if the fonts of both character items are OPENTYPE fonts (as opposed to legacy \TeX fonts), then $\text{LUA}\TeX$ will use the OPENTYPE math algorithm for deciding on the horizontal placement of the super- or subscript.

This works as follows:



- ▶ The vertical position of the script is calculated.
- ▶ The default horizontal position is flat next to the base character.
- ▶ For superscripts, the italic correction of the base character is added.
- ▶ For a superscript, two vertical values are calculated: the bottom of the script (after shifting up), and the top of the base. For a subscript, the two values are the top of the (shifted down) script, and the bottom of the base.
- ▶ For each of these two locations:
 - find the math kern value at this height for the base (for a subscript placement, this is the bottom_right corner, for a superscript placement the top_right corner)
 - find the math kern value at this height for the script (for a subscript placement, this is the top_left corner, for a superscript placement the bottom_left corner)
 - add the found values together to get a preliminary result.
- ▶ The horizontal kern to be applied is the smallest of the two results from previous step.

The math kern value at a specific height is the kern value that is specified by the next higher height and kern pair, or the highest one in the character (if there is no value high enough in the character), or simply zero (if the character has no math kern pairs at all).

7.6.5 Scripts on extensibles: `\Underdelimit`, `\Overdelimit`, `\Udelimitover`, `\Udelimiterunder` and `\Uhexensible`

The primitives `\Underdelimit` and `\Overdelimit` allow the placement of a subscript or superscript on an automatically extensible item and `\Udelimitover` and `\Udelimiterunder` allow the placement of an automatically extensible item as a subscript or superscript on a nucleus. The input:

```
\Uoverdelimit 0 "2194 {\hbox{\strut overdelimit}}$
\Uunderdelimit 0 "2194 {\hbox{\strut underdelimit}}$
\Udelimitover 0 "2194 {\hbox{\strut delimiterover}}$
\Udelimiterunder 0 "2194 {\hbox{\strut delimiterunder}}$
```

will render this:

The vertical placements are controlled by `\Umathunderdelimitergap`, `\Umathunderdelimitervgap`, `\Umathoverdelimitergap`, and `\Umathoverdelimitervgap` in a similar way as limit placements on large operators. The superscript in `\Uoverdelimit` is typeset in a suitable scripted style, the subscript in `\Uunderdelimit` is cramped as well.


These primitives accept an option width specification. When used the also optional keywords `left`, `middle` and `right` will determine what happens when a requested size can't be met (which can happen when we step to successive larger variants).

An extra primitive `\Uhexensible` is available that can be used like this:

```
\Uhexensible width 10cm 0 "2194$
```



This will render this:



Here you can also pass options, like:

```
$\Uhexensible width 1pt middle 0 "2194$
```

This gives:



L^AT_EX internally uses a structure that supports OPENTYPE ‘MathVariants’ as well as TFM ‘extensible recipes’. In most cases where font metrics are involved we have a different code path for traditional fonts and OPENTYPE fonts.

7.6.6 Fractions and the new `\Uskewed` and `\Uskewedwithdelims`

The `\abovewithdelims` command accepts a keyword `exact`. When issued the extra space relative to the rule thickness is not added. One can of course use the `\Umathfraction..gap` commands to influence the spacing. Also the rule is still positioned around the math axis.

```
$$ { {a} \abovewithdelims() exact 4pt {b} }$$
```

The math parameter table contains some parameters that specify a horizontal and vertical gap for skewed fractions. Of course some guessing is needed in order to implement something that uses them. And so we now provide a primitive similar to the other fraction related ones but with a few options so that one can influence the rendering. Of course a user can also mess around a bit with the parameters `\Umathskewedfractionhgap` and `\Umathskewedfractionvgap`.

The syntax used here is:

```
{ {1} \Uskewed / <options> {2} }
{ {1} \Uskewedwithdelims / ( ) <options> {2} }
```

where the options can be `noaxis` and `exact`. By default we add half the axis to the shifts and by default we zero the width of the middle character. For Latin Modern the result looks as follows:

	$x + \frac{a}{b} + x$	$x + \frac{1}{2} + x$	$x + \left(\frac{a}{b}\right) + x$	$x + \left(\frac{1}{2}\right) + x$
<code>exact</code>	$x + \frac{a}{b} + x$	$x + \frac{1}{2} + x$	$x + \left(\frac{a}{b}\right) + x$	$x + \left(\frac{1}{2}\right) + x$
<code>noaxis</code>	$x + \frac{a}{b} + x$	$x + \frac{1}{2} + x$	$x + \left(\frac{a}{b}\right) + x$	$x + \left(\frac{1}{2}\right) + x$
<code>exact noaxis</code>	$x + \frac{a}{b} + x$	$x + \frac{1}{2} + x$	$x + \left(\frac{a}{b}\right) + x$	$x + \left(\frac{1}{2}\right) + x$

7.6.7 Delimiters: `\Uleft`, `\Umiddle` and `\Uright`


Normally you will force delimiters to certain sizes by putting an empty box or rule next to it. The resulting delimiter will either be a character from the stepwise size range or an extensible. The latter can be quite differently positioned than the characters as it depends on the fit as well as the fact if the used characters in the font have depth or height. Commands like (plain T_EXs) `\big` need use this feature. In L^AT_EX we provide a bit more control by three variants that support optional parameters `height`, `depth` and `axis`. The following example uses this:




```

\Uleft   height 30pt depth 10pt      \Udelimiter "0 "0 "000028
\quad x\quad
\Umiddle height 40pt depth 15pt      \Udelimiter "0 "0 "002016
\quad x\quad
\Uright  height 30pt depth 10pt      \Udelimiter "0 "0 "000029
\quad \quad \quad
\Uleft   height 30pt depth 10pt axis \Udelimiter "0 "0 "000028
\quad x\quad
\Umiddle height 40pt depth 15pt axis \Udelimiter "0 "0 "002016
\quad x\quad
\Uright  height 30pt depth 10pt axis \Udelimiter "0 "0 "000029

```



The keyword `exact` can be used as directive that the real dimensions should be applied when the criteria can't be met which can happen when we're still stepping through the successively larger variants. When no dimensions are given the `noaxis` command can be used to prevent shifting over the axis.

You can influence the final class with the keyword `class` which will influence the spacing. The numbers are the same as for character classes.

7.6.8 Accents: `\mathlimitsmode`

When you use `\limits` or `\nolimits` without scripts spacing might get messed up. This can be prevented by setting `\mathlimitsmode` to a non-zero value.

7.7 Extracting values

7.7.1 Codes and using `\Umathcode`, `\Umathcharclass`, `\Umathcharfam` and `\Umathcharslot`

You can extract the components of a math character. Say that we have defined:

```
\Umathcode 1 2 3 4
```

then

```
[\Umathcharclass1] [\Umathcharfam1] [\Umathcharslot1]
```

will return:

```
[2] [3] [4]
```

These commands are provided as convenience. Before they come available you could do the following:



```

\def\Umathcharclass{\directlua{tex.print(tex.getmathcode(token.scan_int())[1])}}
\def\Umathcharfam  {\directlua{tex.print(tex.getmathcode(token.scan_int())[2])}}
\def\Umathcharslot {\directlua{tex.print(tex.getmathcode(token.scan_int())[3])}}

```

7.7.2 Last lines and `\predisplaygapfactor`

There is a new primitive to control the overshoot in the calculation of the previous line in mid-paragraph display math. The default value is 2 times the em width of the current font:

```
\predisplaygapfactor=2000
```

If you want to have the length of the last line independent of math i.e. you don't want to revert to a hack where you insert a fake display math formula in order to get the length of the last line, the following will often work too:

```

\def\lastlinelength{\dimexpr
  \directlua {tex.sprint (
    (nodes.dimensions(node.tail(tex.lists.page_head).list))
  )}sp
\relax}

```

7.8 Math mode

7.8.1 Verbose versions of single-character math commands like `\Usuperscript` and `\Usubscript`

L^AT_EX defines six new primitives that have the same function as `^`, `_`, `$`, and `$$`:

PRIMITIVE	EXPLANATION
<code>\Usuperscript</code>	duplicates the functionality of <code>^</code>
<code>\Usubscript</code>	duplicates the functionality of <code>_</code>
<code>\Ustartmath</code>	duplicates the functionality of <code>\$</code> , when used in non-math mode.
<code>\Ustopmath</code>	duplicates the functionality of <code>\$</code> , when used in inline math mode.
<code>\Ustartdisplaymath</code>	duplicates the functionality of <code>\$\$</code> , when used in non-math mode.
<code>\Ustopdisplaymath</code>	duplicates the functionality of <code>\$\$</code> , when used in display math mode.

The `\Ustopmath` and `\Ustopdisplaymath` primitives check if the current math mode is the correct one (inline vs. displayed), but you can freely intermix the four `mathon/mathoff` commands with explicit dollar sign(s).

7.8.2 Script commands `\Unosuperscript` and `\Unosubscript`

These two commands result in super- and subscripts but with the current style (at the time of rendering). So,

```
\Uleft height 30pt depth 10pt \Udelimiter "0 "0 "000028
```



```

\quad x\quad
\Umiddle height 40pt depth 15pt \Udelimiter "0 "0 "002016
\quad x\quad
\Uright height 30pt depth 10pt \Udelimiter "0 "0 "000029
\quad \quad \quad
\Uleft height 30pt depth 10pt axis \Udelimiter "0 "0 "000028
\quad x\quad
\Umiddle height 40pt depth 15pt axis \Udelimiter "0 "0 "002016
\quad x\quad
\Uright height 30pt depth 10pt axis \Udelimiter "0 "0 "000029

```

results in $x_2^1 = x_2^1 = x_2^1 = x_2^1$.

7.8.3 Allowed math commands in non-math modes

The commands `\mathchar`, and `\Umathchar` and control sequences that are the result of `\mathchardef` or `\Umathchardef` are also acceptable in the horizontal and vertical modes. In those cases, the `\textfont` from the requested math family is used.

7.9 Goodies

7.9.1 Flattening: `\mathflattenmode`

The T_EX math engine collapses ord noads without sub- and superscripts and a character as nucleus. and which has the side effect that in OPENTYPE mode italic corrections are applied (given that they are enabled).

```

\switchtobodyfont[modern]
$V \mathbin{\mathbin{v}} V$\par
$V \mathord{\mathord{v}} V$\par

```

This renders as:

VvV

VvV

When we set `\mathflattenmode` to 31 we get:

$V\,v\,V$

VvV

When you see no difference, then the font probably has the proper character dimensions and no italic correction is needed. For Latin Modern (at least till 2018) there was a visual difference. In that respect this parameter is not always needed unless of course you want efficient math lists anyway.

You can influence flattening by adding the appropriate number to the value of the mode parameter. The default value is 1.



MODE	CLASS
1	ord
2	bin
4	rel
8	punct
16	inner

7.9.2 Less Tracing

Because there are quite some math related parameters and values, it is possible to limit tracing. Only when `tracingassigns` and/or `tracingrestores` are set to 2 or more they will be traced.



8 Nodes

8.1 LUA node representation

TeX's nodes are represented in LUA as userdata objects with a variable set of fields or by a numeric identifier when requested. When you print a node userdata object you will see these numbers. In the following syntax tables, such as the type of such a userdata object is represented as `(node)`.

The return values of `node.types()` are: hlist (0), vlist (1), rule (2), ins (3), mark (4), adjust (5), boundary (6), disc (7), whatsit (8), local_par (9), dir (10), math (11), glue (12), kern (13), penalty (14), style (15), choice (16), parameter (17), noad (18), radical (19), fraction (20), accent (21), fence (22), math_char (23), math_text_char (24), sub_box (25), sub_mlist (26), delim (27), margin_kern (28), glyph (29), unset (30), attribute_list (35), attribute (36), glue_spec (37), temp (38), page_insert (39) and split_insert (40)

In ε -TeX the `\lastnodetype` primitive has been introduced. With this primitive the valid range of numbers is still $[-1, 15]$ and glyph nodes (formerly known as char nodes) have number 0. That way macro packages can use the same symbolic names as in traditional ε -TeX. But you need to keep in mind that these ε -TeX node numbers are different from the real internal ones. When you set `\internalcodesmode` to a non-zero value, the internal codes will be used in the ε -TeX introspection commands `\lastnodetype` and `\currentifttype`.

You can ask for a list of fields with `node.fields` and for valid subtypes with `node.subtypes`. The `node.values` function reports some used values. Valid arguments are `glue`, `style` and `math`. Keep in mind that the setters normally expect a number, but this helper gives you a list of what numbers matter. For practical reason the `pagestate` values are also reported with this helper, but they are backend specific.

The return values of `node.values("glue")` are: normal (0), fi (1), fil (2), fill (3) and filll (4)

The return values of `node.values("style")` are: display (0), crampeddisplay (1), text (2), crampedtext (3), script (4), crampedscript (5), scriptscript (6) and crampedscriptscript (7)

The return values of `node.values("math")` are: quad (0), axis (1), operatorsize (2), overbarkern (3), overbarrule (4), overbarvgap (5), underbarkern (6), underbarrule (7), underbarvgap (8), radicalkern (9), radicalrule (10), radicalvgap (11), radicaldegreebefore (12), radicaldegreeafter (13), radicaldegreeraise (14), stackvgap (15), stacknumup (16), stackdenomdown (17), fractionrule (18), fractionnumvgap (19), fractionnumup (20), fractiondenomvgap (21), fractiondenomdown (22), fractiondelsize (23), skewedfractionhgap (24), skewedfractionvgap (25), limitabovevgap (26), limitabovebgap (27), limitabovekern (28), limitbelowvgap (29), limitbelowbgap (30), limitbelowkern (31), nolimitsubfactor (32), nolimitsupfactor (33), underdelimitervgap (34), underdelimeterbgap (35), overdelimitervgap (36), overdelimeterbgap (37), subshiftdrop (38), supshiftdrop (39), subshiftdown (40), subsupshiftdown (41), subtopmax (42), supshiftdown (43), subbottommin



(44), supsubbottommax (45), subsupvgap (46), spaceafterscript (47), connectoroverlap-min (48), ordordspacing (49), ordopspacing (50), ordbinspacing (51), ordrelspacing (52), ordopenspacing (53), ordclosespacing (54), ordpunctspacing (55), ordinnerspacing (56), opordspacing (57), opopspacing (58), opbinspacing (59), oprelspacing (60), opopenspacing (61), opclosespacing (62), oppunctspacing (63), opinnerspacing (64), binordspacing (65), binopspacing (66), binbinspacing (67), binrelspacing (68), binopenspacing (69), binclosespacing (70), binpunctspacing (71), bininnerspacing (72), relordspacing (73), relopspacing (74), relbinspacing (75), relrelspacing (76), relopenspacing (77), relclosespacing (78), relpunctspacing (79), relinnerspacing (80), openordspacing (81), openopspacing (82), openbinspacing (83), openrelspacing (84), openopenspacing (85), openclosespacing (86), openpunctspacing (87), openinnerspacing (88), closeordspacing (89), closeopspacing (90), closebinspacing (91), closerelspacing (92), closeopenspacing (93), closeclosespacing (94), closepunctspacing (95), closeinnerspacing (96), punctordspacing (97), punctopspacing (98), punctbinspacing (99), punctrelspacing (100), punctopenspacing (101), punctclosespacing (102), punctpunctspacing (103), punctinnerspacing (104), innerordspacing (105), inneropspacing (106), innerbinspacing (107), innerrelspacing (108), inneropenspacing (109), innerclosespacing (110), innerpunctspacing (111) and innerinnerspacing (112)

The return values of `node.values("pagestate"` are: empty (0), box_there (1) and inserts_only (2)

8.2 Main text nodes

These are the nodes that comprise actual typesetting commands. A few fields are present in all nodes regardless of their type, these are:

FIELD	TYPE	EXPLANATION
next	node	the next node in a list, or nil
id	number	the node's type (id) number
subtype	number	the node subtype identifier

The subtype is sometimes just a dummy entry because not all nodes actually use the subtype, but this way you can be sure that all nodes accept it as a valid field name, and that is often handy in node list traversal. In the following tables next and id are not explicitly mentioned.

Besides these three fields, almost all nodes also have an attr field, and there is also a field called prev. That last field is always present, but only initialized on explicit request: when the function `node.slide` is called, it will set up the prev fields to be a backwards pointer in the argument node list. By now most of \TeX 's node processing makes sure that the prev nodes are valid but there can be exceptions, especially when the internal magic uses a leading temp nodes to temporarily store a state.

The \LaTeX engine provides a lot of freedom and it is up to the user to make sure that the node lists remain sane. There are some safeguards but there can be cases where the engine just quits out of frustration. And, of course you can make the engine crash.



8.2.1 hlist and vlist nodes

These lists share fields and subtypes although some subtypes can only occur in horizontal lists while others are unique for vertical lists. The possible fields are `attr`, `depth`, `direction`, `doffset`, `glue_order`, `glue_set`, `glue_sign`, `height`, `hoffset`, `list`, `orientation`, `shift`, `width`, `woffset`, `xoffset` and `yoffset`.

FIELD	TYPE	EXPLANATION
subtype	number	accent, alignment, box, cell, degree, denominator, equation, equationnumber, fraction, hdelimiter, hextensible, indent, limits, line, math, mathchar, nucleus, numerator, over, overdelimiter, radical, scripts, sub, sup, under, underdelimiter, unknown, vdelimiter and vextensible
attr	node	list of attributes
width	number	the width of the box
height	number	the height of the box
depth	number	the depth of the box
direction	number	the direction of this box, see 8.2.15
shift	number	a displacement perpendicular to the character (hlist) or line (vlist) progression direction
glue_order	number	a number in the range [0, 4], indicating the glue order
glue_set	number	the calculated glue ratio
glue_sign	number	0 = normal, 1 = stretching, 2 = shrinking
list	node	the first node of the body of this list

The `orientation`, `woffset`, `hoffset`, `doffset`, `xoffset` and `yoffset` fields are special. They can be used to make the backend rotate and shift boxes which can be handy in for instance vertical typesetting. Because they relate to (and depend on the) the backend they are not discussed here (yet).

A warning: never assign a node list to the `list` field unless you are sure its internal link structure is correct, otherwise an error may result.

Note: the field name `head` and `list` are both valid. Sometimes it makes more sense to refer to a list by `head`, sometimes `list` makes more sense.

8.2.2 rule nodes

Contrary to traditional \TeX , \LaTeX has more `\rule` subtypes because we also use rules to store reusable objects and images. User nodes are invisible and can be intercepted by a callback. The supported fields are `attr`, `data`, `depth`, `height`, `left`, `right`, `width`, `xoffset` and `yoffset`.

FIELD	TYPE	EXPLANATION
subtype	number	box, empty, fraction, image, normal, outline, over, radical, under and user
attr	node	list of attributes
width	number	the width of the rule where the special value <code>-1073741824</code> is used for ‘running’ glue dimensions



height	number	the height of the rule (can be negative)
depth	number	the depth of the rule (can be negative)
left	number	shift at the left end (also subtracted from width)
right	number	(subtracted from width)
dir	string	the direction of this rule, see 8.2.15
index	number	an optional index that can be referred to
transform	number	an private variable (also used to specify outline width)

The left and type right keys are somewhat special (and experimental). When rules are auto adapting to the surrounding box width you can enforce a shift to the right by setting left. The value is also subtracted from the width which can be a value set by the engine itself and is not entirely under user control. The right is also subtracted from the width. It all happens in the backend so these are not affecting the calculations in the frontend (actually the auto settings also happen in the backend). For a vertical rule left affects the height and right affects the depth. There is no matching interface at the \TeX end (although we can have more keywords for rules it would complicate matters and introduce a speed penalty.) However, you can just construct a rule node with LUA and write it to the \TeX input. The outline subtype is just a convenient variant and the transform field specifies the width of the outline.

The xoffset and yoffset fields are special. They can be used to shift rules. Because they relate to (and depend on the) the backend they are not discussed here (yet).

8.2.3 ins nodes

This node relates to the $\backslash\text{insert}$ primitive and support the fields: attr, cost, depth, height, list and spec.

FIELD	TYPE	EXPLANATION
subtype	number	the insertion class
attr	node	list of attributes
cost	number	the penalty associated with this insert
height	number	height of the insert
depth	number	depth of the insert
list	node	the first node of the body of this insert

There is a set of extra fields that concern the associated glue: width, stretch, stretch_order, shrink and shrink_order. These are all numbers.

A warning: never assign a node list to the head field unless you are sure its internal link structure is correct, otherwise an error may result. You can use list instead (often in functions you want to use local variable with similar names and both names are equally sensible).

8.2.4 mark nodes

This one relates to the $\backslash\text{mark}$ primitive and only has a few fields: attr, class and mark.

FIELD	TYPE	EXPLANATION
subtype	number	unused



attr	node	list of attributes
class	number	the mark class
mark	table	a table representing a token list

8.2.5 adjust nodes

This node comes from `\vadjust` primitive and has fields: `attr` and `list`.

FIELD	TYPE	EXPLANATION
subtype	number	normal and pre
attr	node	list of attributes
list	node	adjusted material

A warning: never assign a node list to the head field unless you are sure its internal link structure is correct, otherwise an error may be the result.

8.2.6 disc nodes

The `\discretionary` and `\-`, the `-` character but also the hyphenation mechanism produces these nodes. The available fields are: `attr`, `penalty`, `post`, `pre` and `replace`.

FIELD	TYPE	EXPLANATION
subtype	number	automatic, discretionary, explicit, first, regular and second
attr	node	list of attributes
pre	node	pointer to the pre-break text
post	node	pointer to the post-break text
replace	node	pointer to the no-break text
penalty	number	the penalty associated with the break, normally <code>\hyphenpenalty</code> or <code>\exhyphenpenalty</code>

The subtype numbers 4 and 5 belong to the ‘of-f-ice’ explanation given elsewhere. These disc nodes are kind of special as at some point they also keep information about breakpoints and nested ligatures.

The `pre`, `post` and `replace` fields at the LUA end are in fact indirectly accessed and have a `prev` pointer that is not `nil`. This means that when you mess around with the head of these (three) lists, you also need to reassign them because that will restore the proper `prev` pointer, so:

```
pre = d.pre
-- change the list starting with pre
d.pre = pre
```

Otherwise you can end up with an invalid internal perception of reality and `LUAMETATEX` might even decide to crash on you. It also means that running forward over for instance `pre` is ok but backward you need to stop at `pre`. And you definitely must not mess with the node that `prev` points to, if only because it is not really a node but part of the disc data structure (so freeing it again might crash `LUAMETATEX`).



8.2.7 math nodes

Math nodes represent the boundaries of a math formula, normally wrapped into \$ signs. The following fields are available: attr, shrink, shrink_order, stretch, stretch_order, surround and width.

FIELD	TYPE	EXPLANATION
subtype	number	beginmath and endmath
attr	node	list of attributes
surround	number	width of the \mathsurround kern
width	number	the horizontal or vertical displacement
stretch	number	extra (positive) displacement or stretch amount
stretch_order	number	factor applied to stretch amount
shrink	number	extra (negative) displacement or shrink amount
shrink_order	number	factor applied to shrink amount

The glue fields only kick in when the surround fields is zero.

8.2.8 glue nodes

Skips are about the only type of data objects in traditional T_EX that are not a simple value. They are inserted when T_EX sees a space in the text flow but also by \hskip and \vskip. The structure that represents the glue components of a skip internally is called a glue_spec. In LUAMETAT_EX we don't use the spec itself but just its values. A glue node has the fields: attr, font, leader, shrink, shrink_order, stretch, stretch_order and width.

FIELD	TYPE	EXPLANATION
subtype	number	abovedisplayshortskip, abovedisplayskip, baselineskip, belowdisplayshortskip, belowdisplayskip, cleaders, conditionalmathskip, correctionskip, gleaders, indentskip, intermathskip, leaders, lefthangskip, leftskip, lineskip, mathskip, medmuskip, muglue, parfillskip, parskip, righthangskip, rightskip, spaceskip, splittopskip, tabskip, thickmuskip, thinmuskip, topskip, userskip, xleaders and xspaceskip
attr	node	list of attributes
leader	node	pointer to a box or rule for leaders
width	number	the horizontal or vertical displacement
stretch	number	extra (positive) displacement or stretch amount
stretch_order	number	factor applied to stretch amount
shrink	number	extra (negative) displacement or shrink amount
shrink_order	number	factor applied to shrink amount

Note that we use the key width in both horizontal and vertical glue. This suits the T_EX internals well so we decided to stick to that naming.

The effective width of some glue subtypes depends on the stretch or shrink needed to make the encapsulating box fit its dimensions. For instance, in a paragraph lines normally have glue



representing spaces and these stretch or shrink to make the content fit in the available space. The `effective_glue` function that takes a glue node and a parent (hlist or vlist) returns the effective width of that glue item. When you pass `true` as third argument the value will be rounded.

8.2.9 glue_spec nodes

Internally `LUAMETATEX` (like its ancestors) also uses nodes to store data that is not seen in node lists. For instance the state of expression scanning (`\dimexpr` etc.) and conditionals (`\ifcase` etc.) is also kept in lists of nodes. A glue, which has five components, is stored in a node as well, so, where most registers store just a number, a skip register (of internal quantity) uses a pointer to a glue spec node. It has similar fields as glue nodes: `shrink`, `shrink_order`, `stretch`, `stretch_order` and `width`, which is not surprising because in the past (and other engines than `LUATEX`) a glue node also has its values stored in a glue spec. This has some advantages because often the values are the same, so for instance spacing related skips were not resolved immediately but pointed to the current value of a space related internal register (like `\spaceskip`). But, in `LUATEX` we do resolve these quantities immediately and we put the current values in the glue nodes.

FIELD	TYPE	EXPLANATION
<code>width</code>	number	the horizontal or vertical displacement
<code>stretch</code>	number	extra (positive) displacement or stretch amount
<code>stretch_order</code>	number	factor applied to stretch amount
<code>shrink</code>	number	extra (negative) displacement or shrink amount
<code>shrink_order</code>	number	factor applied to shrink amount

You will only find these nodes in a few places, for instance when you query an internal quantity. In principle we could do without them as we have interfaces that use the five numbers instead. For compatibility reasons we keep glue spec nodes exposed but this might change in the future.

8.2.10 kern nodes

The `\kern` command creates such nodes but for instance the font and math machinery can also add them. There are not that many fields: `attr`, `expansion_factor` and `kern`.

FIELD	TYPE	EXPLANATION
<code>subtype</code>	number	accentkern, fontkern, italiccorrection and userkern
<code>attr</code>	node	list of attributes
<code>kern</code>	number	fixed horizontal or vertical advance
<code>expansion_factor</code>	number	multiplier related to hz for font kerns

8.2.11 penalty nodes

The `\penalty` command is one that generates these nodes. It is one of the type of nodes often found in vertical lists. It has the fields: `attr` and `penalty`.



FIELD	TYPE	EXPLANATION
subtype	number	afterdisplaypenalty, beforesdisplaypenalty, equationnumberpenalty, finalpenalty, linebreakpenalty, linepenalty, noadpenalty, userpenalty and wordpenalty
attr	node	list of attributes
penalty	number	the penalty value

The subtypes are just informative and \TeX itself doesn't use them. When you run into an `linebreakpenalty` you need to keep in mind that it's a accumulation of `club`, `widow` and other relevant penalties.

8.2.12 glyph nodes

These are probably the mostly used nodes and although you can push them in the current list with for instance `\char` \TeX will normally do it for you when it considers some input to be text. Glyph nodes are relatively large and have many fields: `attr`, `char`, `components`, `data`, `depth`, `expansion_factor`, `font`, `height`, `lang`, `left`, `right`, `uchyph`, `width`, `xoffset` and `yoffset`.

FIELD	TYPE	EXPLANATION
subtype	number	bit field
attr	node	list of attributes
char	number	the character index in the font
font	number	the font identifier
lang	number	the language identifier
left	number	the frozen <code>\lefthyphenmn</code> value
right	number	the frozen <code>\righthyphenmn</code> value
uchyph	boolean	the frozen <code>\uchyph</code> value
components	node	pointer to ligature components
xoffset	number	a virtual displacement in horizontal direction
yoffset	number	a virtual displacement in vertical direction
width	number	the (original) width of the character
height	number	the (original) height of the character
depth	number	the (original) depth of the character
expansion_factor	number	the to be applied <code>expansion_factor</code>
data	number	a general purpose field for users (we had room for it)

The `width`, `height` and `depth` values are read-only. The `expansion_factor` is assigned in the `par` builder and used in the backend.

A warning: never assign a node list to the `components` field unless you are sure its internal link structure is correct, otherwise an error may be result. Valid bits for the `subtype` field are:

BIT	MEANING
0	character
1	ligature
2	ghost



- 3 left
- 4 right

The `expansion_factor` has been introduced as part of the separation between font- and backend. It is the result of extensive experiments with a more efficient implementation of expansion. Early versions of L^AT_EX already replaced multiple instances of fonts in the backend by scaling but contrary to P^DF_TE_X in L^AT_EX we now also got rid of font copies in the frontend and replaced them by expansion factors that travel with glyph nodes. Apart from a cleaner approach this is also a step towards a better separation between front- and backend.

The `is_char` function checks if a node is a glyph node with a subtype still less than 256. This function can be used to determine if applying font logic to a glyph node makes sense. The value `nil` gets returned when the node is not a glyph, a character number is returned if the node is still tagged as character and `false` gets returned otherwise. When `nil` is returned, the `id` is also returned. The `is_glyph` variant doesn't check for a subtype being less than 256, so it returns either the character value or `nil` plus the `id`. These helpers are not always faster than separate calls but they sometimes permit making more readable tests. The `uses_font` helpers takes a node and font `id` and returns `true` when a glyph or disc node references that font.

8.2.13 boundary nodes

This node relates to the `\noboundary`, `\boundary`, `\protrusionboundary` and `\wordboundary` primitives. These are small nodes: `attr` and `value` are the only fields.

FIELD	TYPE	EXPLANATION
<code>subtype</code>	number	cancel, protrusion, user and word
<code>attr</code>	node	list of attributes
<code>value</code>	number	values 0-255 are reserved

8.2.14 local_par nodes

This node is inserted at the start of a paragraph. You should not mess too much with this one. Valid fields are: `attr`, `box_left`, `box_left_width`, `box_right`, `box_right_width`, `dir`, `pen_broken` and `pen_inter`.

FIELD	TYPE	EXPLANATION
<code>attr</code>	node	list of attributes
<code>pen_inter</code>	number	local interline penalty (from <code>\localinterlinepenalty</code>)
<code>pen_broken</code>	number	local broken penalty (from <code>\localbrokenpenalty</code>)
<code>dir</code>	string	the direction of this par. see 8.2.15
<code>box_left</code>	node	the <code>\localleftbox</code>
<code>box_left_width</code>	number	width of the <code>\localleftbox</code>
<code>box_right</code>	node	the <code>\localrightbox</code>
<code>box_right_width</code>	number	width of the <code>\localrightbox</code>

A warning: never assign a node list to the `box_left` or `box_right` field unless you are sure its internal link structure is correct, otherwise an error may result.



8.2.15 dir nodes

Direction nodes mark parts of the running text that need a change of direction and the `\textdir` command generates them. Again this is a small node, we just have `attr`, `dir` and `level`.

FIELD	TYPE	EXPLANATION
subtype	number	cancel and normal
attr	node	list of attributes
dir	string	the direction (0 = l2r, 1 = r2l)
level	number	nesting level of this direction

There are only two directions: left-to-right (0) and right-to-left (1). This is different from L^AT_EX that has four directions.

8.2.16 marginkern nodes

Margin kerns result from protrusion and have: `attr`, `glyph` and `width`.

FIELD	TYPE	EXPLANATION
subtype	number	left and right
attr	node	list of attributes
width	number	the advance of the kern
glyph	node	the glyph to be used

8.2.17 Whatsits

A whatsit node is a real simple one and it only has a subtype. It is even less than a user node (which it actually could be) and uses hardly any memory. What you do with it it entirely up to you: it's is real minimalistic. You can assign a subtype and it has attributes. It is all up to the user how they are handled.

8.2.18 Math noads

8.2.18.1 The concept

These are the so-called 'noad's and the nodes that are specifically associated with math processing. When you enter a formula, T_EX creates a node list with regular nodes and noads. Then it hands over the list the math processing engine. The result of that is a nodelist without noads. Most of the noads contain subnodes so that the list of possible fields is actually quite small. Math formulas are both a linked list and a tree. For instance in $e = mc^2$ there is a linked list $e = m c$ but the c has a superscript branch that itself can be a list with branches.

First, there are the objects (the T_EXbook calls them 'atoms') that are associated with the simple math objects: `ord`, `op`, `bin`, `rel`, `open`, `close`, `punct`, `inner`, `over`, `under`, `vcenter`. These all have the same fields, and they are combined into a single node type with separate subtypes for differentiation: `attr`, `nucleus`, `options`, `sub` and `sup`.



Many object fields in math mode are either simple characters in a specific family or math lists or node lists: `math_char`, `math_text_char`, `sub_box` and `sub_mlist` and `delim`. These are endpoints and therefore the next and prev fields of these subnodes are unused.

Some of the more elaborate nodes have an option field. The values in this bitset are common:

MEANING	BITS
set	0x08
internal	0x00 + 0x08
internal	0x01 + 0x08
axis	0x02 + 0x08
no axis	0x04 + 0x08
exact	0x10 + 0x08
left	0x11 + 0x08
middle	0x12 + 0x08
right	0x14 + 0x08
no subscript	0x21 + 0x08
no superscript	0x22 + 0x08
no script	0x23 + 0x08

8.2.18.2 `math_char` and `math_text_char` subnodes

These are the most common ones, as they represent characters, and they both have the same fields: `attr`, `char` and `fam`.

FIELD	TYPE	EXPLANATION
<code>attr</code>	node	list of attributes
<code>char</code>	number	the character index
<code>fam</code>	number	the family number

The `math_char` is the simplest subnode field, it contains the character and family for a single glyph object. The family eventually resolves on a reference to a font. The `math_text_char` is a special case that you will not normally encounter, it arises temporarily during math list conversion (its sole function is to suppress a following italic correction).

8.2.18.3 `sub_box` and `sub_mlist` subnodes

These two subnode types are used for subsidiary list items. For `sub_box`, the `list` points to a 'normal' vbox or hbox. For `sub_mlist`, the `list` points to a math list that is yet to be converted. Their fields are: `attr` and `head`.

FIELD	TYPE	EXPLANATION
<code>attr</code>	node	list of attributes
<code>list</code>	node	list of nodes

A warning: never assign a node list to the `list` field unless you are sure its internal link structure is correct, otherwise an error is triggered.



8.2.18.4 delim subnodes

There is a fifth subnode type that is used exclusively for delimiter fields. As before, the next and prev fields are unused, but we do have: attr, large_char, large_fam, small_char and small_fam.

FIELD	TYPE	EXPLANATION
attr	node	list of attributes
small_char	number	character index of base character
small_fam	number	family number of base character
large_char	number	character index of next larger character
large_fam	number	family number of next larger character

The fields large_char and large_fam can be zero, in that case the font that is set for the small_fam is expected to provide the large version as an extension to the small_char.

8.2.18.5 simple noad nodes

In these noads, the nucleus, sub and sup fields can branch of. Its fields are: attr, nucleus, options, sub and sup.

FIELD	TYPE	EXPLANATION
subtype	number	bin, close, inner, opdisplaylimits, open, oplimits, opnolimits, ord, ordlimits, over, punct, rel, under and vcenter
attr	node	list of attributes
nucleus	kernel node	base
sub	kernel node	subscript
sup	kernel node	superscript
options	number	bitset of rendering options

8.2.18.6 accent nodes

Accent nodes deal with stuff on top or below a math constructs. They support: accent, attr, bot_accent, fraction, nucleus, overlay_accent, sub, sup and top_accent.

FIELD	TYPE	EXPLANATION
subtype	number	bothflexible, fixedboth, fixedbottom and fixedtop
nucleus	kernel node	base
sub	kernel node	subscript
sup	kernel node	superscript
accent	kernel node	top accent
bot_accent	kernel node	bottom accent
fraction	number	larger step criterium (divided by 1000)



8.2.18.7 style nodes

These nodes are signals to switch to another math style. They are quite simple: `attr` and `style`. Currently the subtype is actually used to store the style but don't rely on that for the future. Fields are: `attr` and `style`.

FIELD	TYPE	EXPLANATION
<code>style</code>	string	contains the style

Valid styles are: `display` (0), `crampeddisplay` (1), `text` (2), `crampedtext` (3), `script` (4), `crampedscript` (5), `scriptscript` (6) and `crampedscriptscript` (7).

8.2.18.8 parameter nodes

These nodes are used to (locally) set math parameters: `list`, `name`, `style` and `value`. Fields are: `list`, `name`, `style` and `value`.

FIELD	TYPE	EXPLANATION
<code>style</code>	string	contains the style
<code>name</code>	string	defines the parameter
<code>value</code>	number	holds the value, in case of a <code>mu</code> glue multiple

8.2.18.9 choice nodes

Of its fields `attr`, `display`, `script`, `scriptscript` and `text` most are lists. Warning: never assign a node list unless you are sure its internal link structure is correct, otherwise an error can occur.

FIELD	TYPE	EXPLANATION
<code>attr</code>	node	list of attributes
<code>display</code>	node	list of display size alternatives
<code>text</code>	node	list of text size alternatives
<code>script</code>	node	list of scriptsize alternatives
<code>scriptscript</code>	node	list of scriptscriptsize alternatives

8.2.18.10 radical nodes

Radical nodes are the most complex as they deal with scripts as well as constructed large symbols. Many fields: `attr`, `degree`, `left`, `nucleus`, `options`, `sub`, `sup` and `width`. Warning: never assign a node list to the `nucleus`, `sub`, `sup`, `left`, or `degree` field unless you are sure its internal link structure is correct, otherwise an error can be triggered.

FIELD	TYPE	EXPLANATION
<code>subtype</code>	number	<code>radical</code> , <code>udelimiterover</code> , <code>udelimiterunder</code> , <code>uhextensible</code> , <code>uoverdelimit</code> , <code>uradical</code> , <code>uroot</code> and <code>uunderdelimit</code>
<code>attr</code>	node	list of attributes



nucleus	kernel node	base
sub	kernel node	subscript
sup	kernel node	superscript
left	delimiter node	
degree	kernel node	only set by \Uroot
width	number	required width
options	number	bitset of rendering options

8.2.18.11 fraction nodes

Fraction nodes are also used for delimited cases, hence the left and right fields among: attr, denom, fam, left, middle, num, options, right and width.

FIELD	TYPE	EXPLANATION
attr	node	list of attributes
width	number	(optional) width of the fraction
num	kernel node	numerator
denom	kernel node	denominator
left	delimiter node	left side symbol
right	delimiter node	right side symbol
middle	delimiter node	middle symbol
options	number	bitset of rendering options

Warning: never assign a node list to the num, or denom field unless you are sure its internal link structure is correct, otherwise an error can result.

8.2.18.12 fence nodes

Fence nodes come in pairs but either one can be a dummy (this period driven empty fence). Fields are: attr, class, delim, depth, height, italic and options. Some of these fields are used by the renderer and might get adapted in the process.

FIELD	TYPE	EXPLANATION
subtype	number	left, middle, no, right and unset
attr	node	list of attributes
delim	delimiter node	delimiter specification
italic	number	italic correction
height	number	required height
depth	number	required depth
options	number	bitset of rendering options
class	number	spacing related class



8.3 The node library

8.3.1 Introduction

The node library provides methods that facilitate dealing with (lists of) nodes and their values. They allow you to create, alter, copy, delete, and insert node, the core objects within the typesetter. Nodes are represented in LUA as userdata. The various parts within a node can be accessed using named fields.

Each node has at least the three fields `next`, `id`, and `subtype`. The other available fields depend on the `id`.

- ▶ The `next` field returns the userdata object for the next node in a linked list of nodes, or `nil`, if there is no next node.
- ▶ The `id` indicates TeX's 'node type'. The field `id` has a numeric value for efficiency reasons, but some of the library functions also accept a string value instead of `id`.
- ▶ The `subtype` is another number. It often gives further information about a node of a particular `id`.

Nodes can be compared to each other, but: you are actually comparing indices into the node memory. This means that equality tests can only be trusted under very limited conditions. It will not work correctly in any situation where one of the two nodes has been freed and/or reallocated: in that case, there will be false positives. The general approach to a node related callback is as follows:

- ▶ Assume that the node list that you get is okay and properly double linked. If for some reason the links are not right, you can apply `node.slide` to the list.
- ▶ When you insert a node, make sure you use a previously removed one, a new one or a copy. Don't simply inject the same node twice.
- ▶ When you remove a node, make sure that when this is permanent, you also free the node or list. When you free a node its components are checked and when they are nodes themselves they are also freed.
- ▶ Although you can fool the system, normally you will trigger an error when you try to copy a nonexistent node, or free an already freed node. There is some overhead involved in this checking but the current compromise is acceptable.
- ▶ When you're done, pass back (if needed) the result. It's your responsibility to make sure that the list is properly linked (you can play safe and again apply `node.slide`. In principle you can put nodes in a list that are not acceptable in the following up actions. Some nodes get ignored, others will trigger an error, and sometimes the engine will just crash.

So, from the above it will be clear then memory management of nodes has to be done explicitly by the user. Nodes are not 'seen' by the LUA garbage collector, so you have to call the node freeing functions yourself when you are no longer in need of a node (list). Nodes form linked lists without reference counting, so you have to be careful that when control returns back to L^AT_EX itself, you have not deleted nodes that are still referenced from a next pointer elsewhere, and that you did not create nodes that are referenced more than once. Normally the setters and getters handle this for you.



A good example are discretionary nodes that themselves have three sublists. Internally they use special pointers, but the user never sees them because when you query them or set fields, this property is hidden and taken care of. You just see a list. But, when you mess with these sub lists it is your responsibility that it only contains nodes that are permitted in a discretionary.

There are statistics available with regards to the allocated node memory, which can be handy for tracing. Normally the amount of used nodes is not that large. Typesetting a page can involve thousands of them but most are freed when the page has been shipped out. Compared to other programs, node memory usage is not that excessive. So, if for some reason your application leaks nodes, if at the end of your run you lost as few hundred it's not a real problem. In fact, if you created boxes and made copies but not flushed them for good reason, your run will for sure end with used nodes and the statistics will mention that. The same is true for attributes and skips (glue spec nodes): keeping the current state involves using nodes.

8.3.2 Housekeeping

8.3.2.1 types

This function returns an array that maps node id numbers to node type strings, providing an overview of the possible top-level id types.

```
<table> t = node.types()
```

when we issue this command, we get a table. The currently visible types are { [0] = "hlist", "vlist", "rule", "ins", "mark", "adjust", "boundary", "disc", "whatsit", "local_par", "dir", "math", "glue", "kern", "penalty", "style", "choice", "parameter", "noad", "radical", "fraction", "accent", "fence", "math_char", "math_text_char", "sub_box", "sub_mlist", "delim", "margin_kern", "glyph", "unset", [35] = "attribute_list", [36] = "attribute", [37] = "glue_spec", [38] = "temp", [39] = "page_insert", [40] = "split_insert", } where the numbers are the internal identifiers. Only those nodes are reported that make sense to users so there can be gaps in the range of numbers.

8.3.2.2 id and type

This converts a single type name to its internal numeric representation.

```
<number> id = node.id(<string> type)
```

The `node.id("glyph")` command returns the number 29 and `node.id("hlist")` returns 0 where the number don't relate to importance or some ordering; they just appear in the order that is handy for the engine. Commands like this are rather optimized so performance should be ok but you can of course always store the id in a LUA number.

The reverse operation is: `node.type` If the argument is a number, then the next function converts an internal numeric representation to an external string representation. Otherwise, it will return the string node if the object represents a node, and `nil` otherwise.

```
<string> type = node.type(<any> n)
```



The `node.type(4)` command returns the string `mark` and `node.id(99)` returns `NIL` because there is no node with that id.

8.3.2.3 fields and has_field

This function returns an indexed table with valid field names for a particular type of node.

```
<table> t = node.fields(<number|string> id)
```

The function accepts a string or number, so `node.fields("glyph")` returns `{ [-1] = "prev", [0] = "next", "id", "subtype", "attr", "char", "font", "lang", "left", "right", "uchyph", "components", "xoffset", "yoffset", "width", "height", "depth", "expansion_factor", "data", }` and `node.fields(12)` gives `{ [-1] = "prev", [0] = "next", "id", "subtype", "attr", "leader", "width", "stretch", "shrink", "stretch_order", "shrink_order", "font", }`.

The `has_field` function returns a boolean that is only true if `n` is actually a node, and it has the field.

```
<boolean> t = node.has_field(<node> n, <string> field)
```

This function probably is not that useful but some nodes don't have a subtype, attr or prev field and this is a way to test for that.

8.3.2.4 is_node

```
<boolean|integer> t = node.is_node(<any> item)
```

This function returns a number (the internal index of the node) if the argument is a userdata object of type `<node>` and false when no node is passed.

8.3.2.5 new

The `new` function creates a new node. All its fields are initialized to either zero or `nil` except for `id` and `subtype`. Instead of numbers you can also use strings (names). If you pass a second argument the subtype will be set too.

```
<node> n = node.new(<number|string> id)
<node> n = node.new(<number|string> id, <number|string> subtype)
```

As already has been mentioned, you are responsible for making sure that nodes created this way are used only once, and are freed when you don't pass them back somehow.

8.3.2.6 free, flush_node and flush_list

The next one frees node `n` from `TEX`'s memory. Be careful: no checks are done on whether this node is still pointed to from a register or some next field: it is up to you to make sure that the internal data structures remain correct. Fields that point to nodes or lists are flushed too. So, when you used their content for something else you need to set them to `nil` first.



```
<node> next = node.free(<node> n)
flush_node(<node> n)
```

The `free` function returns the next field of the freed node, while the `flush_node` alternative returns nothing.

A list starting with node `n` can be flushed from T_EX's memory too. Be careful: no checks are done on whether any of these nodes is still pointed to from a register or some next field: it is up to you to make sure that the internal data structures remain correct.

```
node.flush_list(<node> n)
```

When you free for instance a discretionary node, `flush_list` is applied to the `pre`, `post`, `replace` so you don't need to do that yourself. Assigning them `nil` won't free those lists!

8.3.2.7 copy and copy_list

This creates a deep copy of node `n`, including all nested lists as in the case of a `hlist` or `vlist` node. Only the next field is not copied.

```
<node> m = node.copy(<node> n)
```

A deep copy of the node list that starts at `n` can be created too. If `m` is also given, the copy stops just before node `m`.

```
<node> m = node.copy_list(<node> n)
<node> m = node.copy_list(<node> n, <node> m)
```

Note that you cannot copy attribute lists this way. However, there is normally no need to copy attribute lists as when you do assignments to the `attr` field or make changes to specific attributes, the needed copying and freeing takes place automatically. When you change a value of an attribute *in* a list, it will affect all the nodes that share that list.

8.3.2.8 write

```
node.write(<node> n)
```

This function that will append a node list to T_EX's 'current list'. The node list is not deep-copied! There is no error checking either! You might need to enforce horizontal mode in order for this to work as expected.

8.3.3 Manipulating lists

8.3.3.1 slide

This helper makes sure that the node lists is double linked and returns the found tail node.

```
<node> tail = node.slide(<node> n)
```



After some callbacks automatic sliding takes place. This feature can be turned off with `node.fix_node_lists(false)` but you better make sure then that you don't mess up lists. In most cases `TEX` itself only uses next pointers but your other callbacks might expect proper prev pointers too. Future versions of `LUATEX` can add more checking but this will not influence usage.

8.3.3.2 tail

```
<node> m = node.tail(<node> n)
```

Returns the last node of the node list that starts at `n`.

8.3.3.3 length and count

```
<number> i = node.length(<node> n)
<number> i = node.length(<node> n, <node> m)
```

Returns the number of nodes contained in the node list that starts at `n`. If `m` is also supplied it stops at `m` instead of at the end of the list. The node `m` is not counted.

```
<number> i = node.count(<number> id, <node> n)
<number> i = node.count(<number> id, <node> n, <node> m)
```

Returns the number of nodes contained in the node list that starts at `n` that have a matching `id` field. If `m` is also supplied, counting stops at `m` instead of at the end of the list. The node `m` is not counted. This function also accept string `id`'s.

8.3.3.4 remove

```
<node> head, current, removed =
    node.remove(<node> head, <node> current)
<node> head, current =
    node.remove(<node> head, <node> current, <boolean> true)
```

This function removes the node `current` from the list following `head`. It is your responsibility to make sure it is really part of that list. The return values are the new head and current nodes. The returned current is the node following the current in the calling argument, and is only passed back as a convenience (or `nil`, if there is no such node). The returned head is more important, because if the function is called with `current` equal to `head`, it will be changed. When the third argument is passed, the node is freed.

8.3.3.5 insert_before

```
<node> head, new = node.insert_before(<node> head, <node> current, <node> new)
```

This function inserts the node `new` before `current` into the list following `head`. It is your responsibility to make sure that `current` is really part of that list. The return values are the (potentially mutated) head and the node `new`, set up to be part of the list (with correct next field). If `head` is initially `nil`, it will become `new`.



8.3.3.6 insert_after

```
<node> head, new = node.insert_after(<node> head, <node> current, <node> new)
```

This function inserts the node `new` after `current` into the list following `head`. It is your responsibility to make sure that `current` is really part of that list. The return values are the head and the node `new`, set up to be part of the list (with correct next field). If `head` is initially `nil`, it will become `new`.

8.3.3.7 last_node

```
<node> n = node.last_node()
```

This function pops the last node from T_EX's 'current list'. It returns that node, or `nil` if the current list is empty.

8.3.3.8 traverse

```
<node> t, id, subtype = node.traverse(<node> n)
```

This is a LUA iterator that loops over the node list that starts at `n`. Typically code looks like this:

```
for n in node.traverse(head) do
    ...
end
```

is functionally equivalent to:

```
do
    local n
    local function f (head,var)
        local t
        if var == nil then
            t = head
        else
            t = var.next
        end
        return t
    end
    while true do
        n = f (head, n)
        if n == nil then break end
        ...
    end
end
```

It should be clear from the definition of the function `f` that even though it is possible to add or remove nodes from the node list while traversing, you have to take great care to make sure all the next (and prev) pointers remain valid.



If the above is unclear to you, see the section ‘For Statement’ in the LUA Reference Manual.

8.3.3.9 `traverse_id`

`<node> t, subtype = node.traverse_id(<number> id, <node> n)`

This is an iterator that loops over all the nodes in the list that starts at `n` that have a matching `id` field.

See the previous section for details. The change is in the local function `f`, which now does an extra while loop checking against the upvalue `id`:

```
local function f(head,var)
  local t
  if var == nil then
    t = head
  else
    t = var.next
  end
  while not t.id == id do
    t = t.next
  end
  return t
end
```

8.3.3.10 `traverse_char` and `traverse_glyph`

The `traverse_char` iterator loops over the glyph nodes in a list. Only nodes with a subtype less than 256 are seen.

`<node> n, font, char = node.traverse_char(<node> n)`

The `traverse_glyph` iterator loops over a list and returns the list and filters all glyphs:

`<node> n, font, char = node.traverse_glyph(<node> n)`

8.3.3.11 `traverse_list`

This iterator loops over the `hlist` and `vlist` nodes in a list.

`<node> n, id, subtype, list = node.traverse_list(<node> n)`

The four return values can save some time compared to fetching these fields but in practice you seldom need them all. So consider it a (side effect of experimental) convenience.

8.3.3.12 `find_node`

This helper returns the location of the first match at or after node `n`:



```
<node> n = node.traverse_list(<node> n, <integer> subtype)
<node> n, subtype = node.traverse_list(<node> n)
```

8.3.4 Glue handling

8.3.4.1 setglue

You can set the properties of a glue in one go. If you pass no values, the glue will become a zero glue.

```
node.setglue(<node> n)
node.setglue(<node> n,width,stretch,shrink,stretch_order,shrink_order)
```

When you pass values, only arguments that are numbers are assigned so

```
node.setglue(n,655360,false,65536)
```

will only adapt the width and shrink.

When a list node is passed, you set the glue, order and sign instead.

8.3.4.2 getglue

The next call will return 5 values or nothing when no glue is passed.

```
<integer> width, <integer> stretch, <integer> shrink, <integer> stretch_order,
<integer> shrink_order = node.getglue(<node> n)
```

When the second argument is false, only the width is returned (this is consistent with `tex.get`).

When a list node is passed, you get back the glue that is set, the order of that glue and the sign.

8.3.4.3 is_zero_glue

This function returns true when the width, stretch and shrink properties are zero.

```
<boolean> isglue = node.is_zero_glue(<node> n)
```

8.3.5 Attribute handling

8.3.5.1 Attributes

Assignments to attributes registers result in assigning lists with set attributes to nodes and the implementation is non-trivial because the value that is attached to a node is essentially a (sorted) sparse array of key-value pairs. It is generally easiest to deal with attribute lists and attributes by using the dedicated functions in the node library.



8.3.5.2 attribute_list nodes

An `attribute_list` item is used as a head pointer for a list of attribute items. It has only one user-visible field:

FIELD	TYPE	EXPLANATION
next	node	pointer to the first attribute

8.3.5.3 attr nodes

A normal node's attribute field will point to an item of type `attribute_list`, and the next field in that item will point to the first defined 'attribute' item, whose next will point to the second 'attribute' item, etc.

FIELD	TYPE	EXPLANATION
next	node	pointer to the next attribute
number	number	the attribute type id
value	number	the attribute value

As mentioned it's better to use the official helpers rather than edit these fields directly. For instance the `prev` field is used for other purposes and there is no double linked list.

8.3.5.4 current_attr

This returns the currently active list of attributes, if there is one.

```
<node> m = node.current_attr()
```

The intended usage of `current_attr` is as follows:

```
local x1 = node.new("glyph")
x1.attr = node.current_attr()
local x2 = node.new("glyph")
x2.attr = node.current_attr()
```

or:

```
local x1 = node.new("glyph")
local x2 = node.new("glyph")
local ca = node.current_attr()
x1.attr = ca
x2.attr = ca
```

The attribute lists are ref counted and the assignment takes care of incrementing the refcount. You cannot expect the value `ca` to be valid any more when you assign attributes (using `tex.setattribute`) or when control has been passed back to `TEX`.



8.3.5.5 has_attribute

```
<number> v = node.has_attribute(<node> n, <number> id)
<number> v = node.has_attribute(<node> n, <number> id, <number> val)
```

Tests if a node has the attribute with number `id` set. If `val` is also supplied, also tests if the value matches `val`. It returns the value, or, if no match is found, `nil`.

8.3.5.6 get_attribute

```
<number> v = node.get_attribute(<node> n, <number> id)
```

Tests if a node has an attribute with number `id` set. It returns the value, or, if no match is found, `nil`. If no `id` is given then the zero attributes is assumed.

8.3.5.7 find_attribute

```
<number> v, <node> n = node.find_attribute(<node> n, <number> id)
```

Finds the first node that has attribute with number `id` set. It returns the value and the node if there is a match and otherwise nothing.

8.3.5.8 set_attribute

```
node.set_attribute(<node> n, <number> id, <number> val)
```

Sets the attribute with number `id` to the value `val`. Duplicate assignments are ignored.

8.3.5.9 unset_attribute

```
<number> v =
  node.unset_attribute(<node> n, <number> id)
<number> v =
  node.unset_attribute(<node> n, <number> id, <number> val)
```

Unsets the attribute with number `id`. If `val` is also supplied, it will only perform this operation if the value matches `val`. Missing attributes or attribute-value pairs are ignored.

If the attribute was actually deleted, returns its old value. Otherwise, returns `nil`.

8.3.6 Glyph handling

8.3.6.1 first_glyph

```
<node> n = node.first_glyph(<node> n)
<node> n = node.first_glyph(<node> n, <node> m)
```



Returns the first node in the list starting at `n` that is a glyph node with a subtype indicating it is a glyph, or `nil`. If `m` is given, processing stops at (but including) that node, otherwise processing stops at the end of the list.

8.3.6.2 `is_char` and `is_glyph`

The subtype of a glyph node signals if the glyph is already turned into a character reference or not.

```
<boolean> b = node.is_char(<node> n)
<boolean> b = node.is_glyph(<node> n)
```

8.3.6.3 `has_glyph`

This function returns the first glyph or disc node in the given list:

```
<node> n = node.has_glyph(<node> n)
```

8.3.6.4 `ligaturing`

```
<node> h, <node> t, <boolean> success = node.ligaturing(<node> n)
<node> h, <node> t, <boolean> success = node.ligaturing(<node> n, <node> m)
```

Apply T_EX-style ligaturing to the specified nodelist. The tail node `m` is optional. The two returned nodes `h` and `t` are the new head and tail (both `n` and `m` can change into a new ligature).

8.3.6.5 `kerning`

```
<node> h, <node> t, <boolean> success = node.kerning(<node> n)
<node> h, <node> t, <boolean> success = node.kerning(<node> n, <node> m)
```

Apply T_EX-style kerning to the specified node list. The tail node `m` is optional. The two returned nodes `h` and `t` are the head and tail (either one of these can be an inserted kern node, because special kernings with word boundaries are possible).

8.3.6.6 `unprotect_glyph[s]`

```
node.unprotect_glyph(<node> n)
node.unprotect_glyphs(<node> n, [<node> n])
```

Subtracts 256 from all glyph node subtypes. This and the next function are helpers to convert from characters to glyphs during node processing. The second argument is optional and indicates the end of a range.

8.3.6.7 `protect_glyph[s]`

```
node.protect_glyph(<node> n)
```



```
node.protect_glyphs(<node> n, [<node> n])
```

Adds 256 to all glyph node subtypes in the node list starting at n, except that if the value is 1, it adds only 255. The special handling of 1 means that characters will become glyphs after subtraction of 256. A single character can be marked by the singular call. The second argument is optional and indicates the end of a range.

8.3.6.8 protrusion_skippable

```
<boolean> skippable = node.protrusion_skippable(<node> n)
```

Returns true if, for the purpose of line boundary discovery when character protrusion is active, this node can be skipped.

8.3.6.9 check_discretionary, check_discretionaries

When you fool around with disc nodes you need to be aware of the fact that they have a special internal data structure. As long as you reassign the fields when you have extended the lists it's ok because then the tail pointers get updated, but when you add to list without reassigning you might end up in trouble when the linebreak routine kicks in. You can call this function to check the list for issues with disc nodes.

```
node.check_discretionary(<node> n)
node.check_discretionaries(<node> head)
```

The plural variant runs over all disc nodes in a list, the singular variant checks one node only (it also checks if the node is a disc node).

8.3.6.10 flatten_discretionaries

This function will remove the discretionaries in the list and inject the replace field when set.

```
<node> head, count = node.flatten_discretionaries(<node> n)
```

8.3.7 Packaging

8.3.7.1 hpack

This function creates a new hlist by packaging the list that begins at node n into a horizontal box. With only a single argument, this box is created using the natural width of its components. In the three argument form, info must be either additional or exactly, and w is the additional (\hbox spread) or exact (\hbox to) width to be used. The second return value is the badness of the generated box.

```
<node> h, <number> b =
    node.hpack(<node> n)
```



```

<node> h, <number> b =
    node.hpack(<node> n, <number> w, <string> info)
<node> h, <number> b =
    node.hpack(<node> n, <number> w, <string> info, <string> dir)

```

Caveat: there can be unexpected side-effects to this function, like updating some of the `\marks` and `\inserts`. Also note that the content of `h` is the original node list `n`: if you call `node.free(h)` you will also free the node list itself, unless you explicitly set the `list` field to `nil` beforehand. And in a similar way, calling `node.free(n)` will invalidate `h` as well!

8.3.7.2 vpack

This function creates a new `vlist` by packaging the list that begins at node `n` into a vertical box. With only a single argument, this box is created using the natural height of its components. In the three argument form, `info` must be either `additional` or `exactly`, and `w` is the additional (`\vbox spread`) or exact (`\vbox to`) height to be used.

```

<node> h, <number> b =
    node.vpack(<node> n)
<node> h, <number> b =
    node.vpack(<node> n, <number> w, <string> info)
<node> h, <number> b =
    node.vpack(<node> n, <number> w, <string> info, <string> dir)

```

The second return value is the badness of the generated box. See the description of `hpack` for a few memory allocation caveats.

8.3.7.3 prepend_prevdepth

This function is somewhat special in the sense that it is an experimental helper that adds the interlinespace to a line keeping the `baselineskip` and `lineskip` into account.

```

<node> n, <number> delta =
    node.prepend_prevdepth(<node> n, <number> prevdepth)

```

8.3.7.4 dimensions, rangedimensions, naturalwidth

```

<number> w, <number> h, <number> d =
    node.dimensions(<node> n)
<number> w, <number> h, <number> d =
    node.dimensions(<node> n, <node> t)

```

This function calculates the natural in-line dimensions of the node list starting at node `n` and terminating just before node `t` (or the end of the list, if there is no second argument). The return values are scaled points. An alternative format that starts with glue parameters as the first three arguments is also possible:

```

<number> w, <number> h, <number> d =

```



```

    node.dimensions(<number> glue_set, <number> glue_sign, <number> glue_order,
        <node> n)
<number> w, <number> h, <number> d =
    node.dimensions(<number> glue_set, <number> glue_sign, <number> glue_order,
        <node> n, <node> t)

```

This calling method takes glue settings into account and is especially useful for finding the actual width of a sublist of nodes that are already boxed, for example in code like this, which prints the width of the space in between the a and b as it would be if `\box0` was used as-is:

```

\setbox0 = \hbox to 20pt {a b}

\directlua{print (node.dimensions(
    tex.box[0].glue_set,
    tex.box[0].glue_sign,
    tex.box[0].glue_order,
    tex.box[0].head.next,
    node.tail(tex.box[0].head)
)) }

```

You need to keep in mind that this is one of the few places in \TeX where floats are used, which means that you can get small differences in rounding when you compare the width reported by `hpack` with `dimensions`.

The second alternative saves a few lookups and can be more convenient in some cases:

```

<number> w, <number> h, <number> d =
    node.rangedimensions(<node> parent, <node> first)
<number> w, <number> h, <number> d =
    node.rangedimensions(<node> parent, <node> first, <node> last)

```

A simple and somewhat more efficient variant is this:

```

<number> w =
    node.naturalwidth(<node> start, <node> stop)

```

8.3.8 Math

8.3.8.1 mlist_to_hlist

```

<node> h =
    node.mlist_to_hlist(<node> n, <string> display_type, <boolean> penalties)

```

This runs the internal mlist to hlist conversion, converting the math list in `n` into the horizontal list `h`. The interface is exactly the same as for the callback `mlist_to_hlist`.

8.3.8.2 end_of_math

```

<node> t = node.end_of_math(<node> start)

```



Looks for and returns the next `math_node` following the `start`. If the given node is a math end node this helper returns that node, else it follows the list and returns the next math endnote. If no such node is found `nil` is returned.

8.4 Two access models

Deep down in $\text{T}_{\text{E}}\text{X}$ a node has a number which is a numeric entry in a memory table. In fact, this model, where $\text{T}_{\text{E}}\text{X}$ manages memory is real fast and one of the reasons why plugging in callbacks that operate on nodes is quite fast too. Each node gets a number that is in fact an index in the memory table and that number often is reported when you print node related information. You go from userdata nodes and there numeric references and back with:

```
<integer> d = node.todirect(<node> n)
<node> n = node.tonode(<integer> d)
```

The userdata model is rather robust as it is a virtual interface with some additional checking while the more direct access which uses the node numbers directly. However, even with userdata you can get into troubles when you free nodes that are no longer allocated or mess up lists. if you apply `tostring` to a node you see its internal (direct) number and id.

The first model provides key based access while the second always accesses fields via functions:

```
nodeobject.char
getfield(nodenum, "char")
```

If you use the direct model, even if you know that you deal with numbers, you should not depend on that property but treat it as an abstraction just like traditional nodes. In fact, the fact that we use a simple basic datatype has the penalty that less checking can be done, but less checking is also the reason why it's somewhat faster. An important aspect is that one cannot mix both methods, but you can cast both models. So, multiplying a node number makes no sense.

So our advice is: use the indexed (table) approach when possible and investigate the direct one when speed might be a real issue. For that reason $\text{LUA}_{\text{T}_{\text{E}}\text{X}}$ also provide the `get*` and `set*` functions in the top level node namespace. There is a limited set of getters. When implementing this direct approach the regular index by key variant was also optimized, so direct access only makes sense when nodes are accessed millions of times (which happens in some font processing for instance).

We're talking mostly of getters because setters are less important. Documents have not that many content related nodes and setting many thousands of properties is hardly a burden contrary to millions of consultations.

Normally you will access nodes like this:

```
local next = current.next
if next then
    -- do something
end
```

Here `next` is not a real field, but a virtual one. Accessing it results in a metatable method being called. In practice it boils down to looking up the node type and based on the node type checking



for the field name. In a worst case you have a node type that sits at the end of the lookup list and a field that is last in the lookup chain. However, in successive versions of L^AT_EX these lookups have been optimized and the most frequently accessed nodes and fields have a higher priority.

Because in practice the next accessor results in a function call, there is some overhead involved. The next code does the same and performs a tiny bit faster (but not that much because it is still a function call but one that knows what to look up).

```
local next = node.next(current)
if next then
    -- do something
end
```

In the direct namespace there are more helpers and most of them are accompanied by setters. The getters and setters are clever enough to see what node is meant. We don't deal with whatsit nodes: their fields are always accessed by name. It doesn't make sense to add getters for all fields, we just identifier the most likely candidates. In complex documents, many node and fields types never get seen, or seen only a few times, but for instance glyphs are candidates for such optimization. The `node.direct` interface has some more helpers.²

The `setdisc` helper takes three (optional) arguments plus an optional fourth indicating the sub-type. Its `getdisc` takes an optional boolean; when its value is `true` the tail nodes will also be returned. The `setfont` helper takes an optional second argument, it being the character. The `directmode` setter `setlink` takes a list of nodes and will link them, thereby ignoring `nil` entries. The first valid node is returned (beware: for good reason it assumes single nodes). For rarely used fields no helpers are provided and there are a few that probably are used seldom too but were added for consistency. You can of course always define additional accessors using `get-field` and `setfield` with little overhead. When the second argument of `setattributelist` is `true` the current attribute list is assumed.

In C^ON^TE^XT some of the not performance-critical userdata variants are emulated in L^AU^A and not in the engine, so we retain downward compatibility.

FUNCTION	NODE	DIRECT	emulated
<code>check_discretionaries</code>	–	+	+
<code>check_discretionary</code>	–	+	+
<code>copy</code>	+	+	
<code>copy_list</code>	+	+	
<code>count</code>	–	+	+
<code>current_attr</code>	+	+	
<code>dimensions</code>	–	+	+
<code>effective_glue</code>	–	+	+
<code>end_of_math</code>	–	+	+
<code>find_attribute</code>	–	+	+
<code>first_glyph</code>	–	+	+
<code>flatten_discretionaries</code>	–	+	+

² We can define the helpers in the node namespace with `getfield` which is about as efficient, so at some point we might provide that as module.



flush_list	+	+	
flush_node	+	+	
free	+	+	
get_attribute	+	+	
get_properties_table	+	+	
get_synctex_fields	–	+	
getattributelist	–	+	
getboth	–	+	
getbox	–	+	
getchar	–	+	
getcomponents	–	+	
getdata	–	+	
getdepth	–	+	
getdirection	–	+	
getdisc	–	+	
getexpansion	–	+	
getfam	–	+	
getfield	+	+	
getfont	–	+	
getglue	–	+	+
getheight	–	+	
getid	–	+	
getkern	–	+	
getlang	–	+	
getleader	–	+	
getlist	–	+	
getnext	–	+	
getnormalizedline	–	+	
getnucleus	–	+	
getoffsets	–	+	
getorientation	–	+	
getpenalty	–	+	
getpost	–	+	
getpre	–	+	
getprev	–	+	
getproperty	+	+	
getreplace	–	+	
getshift	–	+	
getsub	–	+	
getsubtype	–	+	
getsup	–	+	
getwhd	–	+	
getwidth	–	+	
has_attribute	+	+	
has_dimensions	–	+	
has_field	+	+	



has_glyph	-	+	+
hpack	-	+	+
insert_after	+	+	
insert_before	+	+	
is_char	-	+	
is_direct	-	+	
is_glyph	-	+	
is_node	+	+	
is_valid	-	+	
is_zero_glue	-	+	+
kerning	-	+	+
last_node	-	+	+
length	-	+	+
ligaturing	-	+	+
make_extensible	-	+	+
mlist_to_hlist	-	+	+
naturalwidth	-	+	+
new	+	+	
prepend_prevdepth	-	+	+
protect_glyphs	-	+	+
protect_glyph	-	+	+
protrusion_skippable	-	+	+
rangedimensions	-	+	+
remove	+	+	
set_attribute	+	+	
set_synctex_fields	-	+	
setattributelist	-	+	
setboth	-	+	
setbox	-	+	
setchar	-	+	
setcomponents	-	+	
setdata	-	+	
setdepth	-	+	
setdirection	-	+	
setdisc	-	+	
setexpansion	-	+	
setfam	-	+	
setfield	+	+	
setfont	-	+	
setglue	+	+	
setheight	-	+	
setkern	-	+	
setlang	-	+	
setleader	-	+	
setlink	-	+	
setlist	-	+	



setnext	-	+	
setnucleus	-	+	
setoffsets	-	+	
setorientation	-	+	
setpenalty	-	+	
setprev	-	+	
setproperty	+	+	
setshift	-	+	
setsplit	-	+	
setsub	-	+	
setsubtype	-	+	
setup	-	+	
setwhd	-	+	
setwidth	-	+	
slide	-	+	+
start_of_par	-	+	
subtype	-	-	
tail	+	+	
todirect	-	+	
tonode	-	+	
tostring	+	-	
traverse	+	+	
traverse_char	+	+	
traverse_glyph	+	+	
traverse_id	+	+	
traverse_list	+	+	
type	+	-	
unprotect_glyphs	-	+	+
unprotect_glyph	-	+	+
unset_attribute	+	+	
usedlist	-	+	+
uses_font	-	+	+
vpack	-	+	+
write	+	+	

The `node.next` and `node.prev` functions will stay but for consistency there are variants called `getnext` and `getprev`. We had to use `get` because `node.id` and `node.subtype` are already taken for providing meta information about nodes. Note: The getters do only basic checking for valid keys. You should just stick to the keys mentioned in the sections that describe node properties.

Some of the getters and setters handle multiple node types, given that the field is relevant. In that case, some field names are considered similar (like `kern` and `width`, or `data` and `value`. In retrospect we could have normalized field names better but we decided to stick to the original (internal) names as much as possible. After all, at the LUA end one can easily create synonyms.

Some nodes have indirect references. For instance a math character refers to a family instead of a font. In that case we provide a virtual font field as accessor. So, `getfont` and `.font` can be



used on them. The same is true for the width, height and depth of glue nodes. These actually access the spec node properties, and here we can set as well as get the values.

In some places L^AT_EX can do a bit of extra checking for valid node lists and you can enable that with:

```
node.fix_node_lists(<boolean> b)
```

You can set and query the SYNCT_EX fields, a file number aka tag and a line number, for a glue, kern, hlist, vlist, rule and math nodes as well as glyph nodes (although this last one is not used in native SYNCT_EX).

```
node.set_synctex_fields(<integer> f, <integer> l)
<integer> f, <integer> l =
    node.get_synctex_fields(<node> n)
```

Of course you need to know what you're doing as no checking on sane values takes place. Also, the synctex interpreter used in editors is rather peculiar and has some assumptions (heuristics).

8.5 Normalization

As an experiment the lines resulting from paragraph construction can be normalized. There are several modes, that can be set and queried with:

```
node.direct.setnormalize(<integer> n)
<integer> n = node.direct.getnormalize()
```

The state of a line (a hlist) can be queried with:

```
<integer> leftskip, <integer> rightskip,
    <integer> lefthangskip, <integer> righthangskip,
    <node> head, <node> tail,
    <integer> parindent, <integer> parfillskip = node.direct.getnormalized()
```

The modes accumulate, so mode 4 includes 1 upto 3:

VALUE	EXPLANATION
1	left and right skips and directions
2	indentation and parfill skip
3	hanging indentation and par shapes
4	idem but before left and right skips
5	inject compensation for overflow

This is experimental code and might take a while to become frozen.

8.6 Properties

Attributes are a convenient way to relate extra information to a node. You can assign them at the T_EX end as well as at the L^AU_A end and consult them at the L^AU_A end. One big advantage is that



they obey grouping. They are linked lists and normally checking for them is pretty efficient, even if you use a lot of them. A macro package has to provide some way to manage these attributes at the T_EX end because otherwise clashes in their usage can occur.

Each node also can have a properties table and you can assign values to this table using the `setproperty` function and get properties using the `getproperty` function. Managing properties is way more demanding than managing attributes.

Take the following example:

```
\directlua {
  local n = node.new("glyph")

  node.setproperty(n,"foo")
  print(node.getproperty(n))

  node.setproperty(n,"bar")
  print(node.getproperty(n))

  node.free(n)
}
```

This will print `foo` and `bar` which in itself is not that useful when multiple mechanisms want to use this feature. A variant is:

```
\directlua {
  local n = node.new("glyph")

  node.setproperty(n,{ one = "foo", two = "bar" })
  print(node.getproperty(n).one)
  print(node.getproperty(n).two)

  node.free(n)
}
```

This time we store two properties with the node. It really makes sense to have a table as property because that way we can store more. But in order for that to work well you need to do it this way:

```
\directlua {
  local n = node.new("glyph")

  local t = node.getproperty(n)

  if not t then
    t = { }
    node.setproperty(n,t)
  end
  t.one = "foo"
```



```

t.two = "bar"

print(node.getproperty(n).one)
print(node.getproperty(n).two)

node.free(n)
}

```

Here our own properties will not overwrite other users properties unless of course they use the same keys. So, eventually you will end up with something:

```

\directlua {
    local n = node.new("glyph")

    local t = node.getproperty(n)

    if not t then
        t = { }
        node.setproperty(n,t)
    end
    t.myself = { one = "foo", two = "bar" }

    print(node.getproperty(n).myself.one)
    print(node.getproperty(n).myself.two)

    node.free(n)
}

```

This assumes that only you use `myself` as subtable. The possibilities are endless but care is needed. For instance, the generic font handler that ships with `CONTEXT` uses the injections subtable and you should not mess with that one!

There are a few helper functions that you normally should not touch as user: `flush_properties_table` will wipe the table (normally a bad idea), `get_properties_table` will give the table that stores properties (using direct entries) and you can best not mess too much with that one either because `LUATEX` itself will make sure that entries related to nodes will get wiped when nodes get freed, so that the LUA garbage collector can do its job. In fact, the main reason why we have this mechanism is that it saves the user (or macro package) some work. One can easily write a property mechanism in LUA where after a shipout properties gets cleaned up but it's not entirely trivial to make sure that with each freed node also its properties get freed, due to the fact that there can be nodes left over for a next page. And having a callback bound to the node deallocator would add way too much overhead.

When we copy a node list that has a table as property, there are several possibilities: we do the same as a new node, we copy the entry to the table in properties (a reference), we do a deep copy of a table in the properties, we create a new table and give it the original one as a metatable. After some experiments (that also included timing) with these scenarios we decided that a deep copy made no sense, nor did nilling. In the end both the shallow copy and the metatable variant were both ok, although the second one is slower. The most important aspect to keep in mind is



that references to other nodes in properties no longer can be valid for that copy. We could use two tables (one unique and one shared) or metatables but that only complicates matters.

When defining a new node, we could already allocate a table but it is rather easy to do that at the lua end e.g. using a metatable `__index` method. That way it is under macro package control. When deleting a node, we could keep the slot (e.g. setting it to false) but it could make memory consumption raise unneeded when we have temporary large node lists and after that only small lists. Both are not done.

So in the end this is what happens now: when a node is copied, and it has a table as property, the new node will share that table. If the second argument of `set_properties_mode` is true then a metatable approach is chosen: the copy gets its own table with the original table as metatable. If you use the generic font loader the mode is enabled that way.

A few more experiments were done. For instance: copy attributes to the properties so that we have fast access at the LUA end. In the end the overhead is not compensated by speed and convenience, in fact, attributes are not that slow when it comes to accessing them. So this was rejected.

Another experiment concerned a bitset in the node but again the gain compared to attributes was neglectable and given the small amount of available bits it also demands a pretty strong agreement over what bit represents what, and this is unlikely to succeed in the T_EX community. It doesn't pay off.

Just in case one wonders why properties make sense: it is not so much speed that we gain, but more convenience: storing all kind of (temporary) data in attributes is no fun and this mechanism makes sure that properties are cleaned up when a node is freed. Also, the advantage of a more or less global properties table is that we stay at the LUA end. An alternative is to store a reference in the node itself but that is complicated by the fact that the register has some limitations (no numeric keys) and we also don't want to mess with it too much.





9 LUA callbacks

9.1 Registering callbacks

This library has functions that register, find and list callbacks. Callbacks are LUA functions that are called in well defined places. There are two kind of callbacks: those that mix with existing functionality, and those that (when enabled) replace functionality. In mostly cases the second category is expected to behave similar to the built in functionality because in a next step specific data is expected. For instance, you can replace the hyphenation routine. The function gets a list that can be hyphenated (or not). The final list should be valid and is (normally) used for constructing a paragraph. Another function can replace the ligature builder and/or kerner. Doing something else is possible but in the end might not give the user the expected outcome.

The first thing you need to do is registering a callback:

```
id, error =  
    callback.register(<string> callback_name, <function> func)  
id, error =  
    callback.register(<string> callback_name, nil)  
id, error =  
    callback.register(<string> callback_name, false)
```

Here the `callback_name` is a predefined callback name, see below. The function returns the internal id of the callback or `nil`, if the callback could not be registered. In the latter case, `error` contains an error message, otherwise it is `nil`.

LUA_T_E_X internalizes the callback function in such a way that it does not matter if you redefine a function accidentally.

Callback assignments are always global. You can use the special value `nil` instead of a function for clearing the callback.

For some minor speed gain, you can assign the boolean `false` to the non-file related callbacks, doing so will prevent LUA_T_E_X from executing whatever it would execute by default (when no callback function is registered at all). Be warned: this may cause all sorts of grief unless you know *exactly* what you are doing!

```
<table> info =  
    callback.list()
```

The keys in the table are the known callback names, the value is a boolean where `true` means that the callback is currently set (active).

```
<function> f = callback.find(callback_name)
```

If the callback is not set, `find` returns `nil`.

9.2 File related callbacks

The behaviour documented in this subsection is considered stable in the sense that there will not be backward-incompatible changes any more.



9.2.1 find_read_file

Your callback function should have the following conventions:

```
<string> actual_name =  
    function (<number> id_number, <string> asked_name)
```

Arguments:

id_number

This number is zero for the log or `\input` files. For `TEX`'s `\read` or `\write` the number is incremented by one, so `\read0` becomes 1.

asked_name

This is the user-supplied filename, as found by `\input`, `\openin` or `\openout`.

Return value:

actual_name

This is the filename used. For the very first file that is read in by `TEX`, you have to make sure you return an `actual_name` that has an extension and that is suitable for use as `jobname`. If you don't, you will have to manually fix the name of the log file and output file after `LUATEX` is finished, and an eventual format filename will become mangled. That is because these file names depend on the `jobname`.

You have to return `nil` if the file cannot be found.

9.2.2 find_data_file

Your callback function should have the following conventions:

```
<string> actual_name =  
    function (<string> asked_name)
```

Return `nil` if the file cannot be found.

9.2.3 find_format_file

Your callback function should have the following conventions:

```
<string> actual_name =  
    function (<string> asked_name)
```

The `asked_name` is a format file for reading (the format file for writing is always opened in the current directory).

9.2.4 open_read_file

Your callback function should have the following conventions:

```
<table> env =
```



```
function (<string> file_name)
```

Argument:

file_name

The filename returned by a previous `find_read_file` or the return value of `kpse.find_file()` if there was no such callback defined.

Return value:

env

This is a table containing at least one required and one optional callback function for this file. The required field is `reader` and the associated function will be called once for each new line to be read, the optional one is `close` that will be called once when `LUATEX` is done with the file.

`LUATEX` never looks at the rest of the table, so you can use it to store your private per-file data. Both the callback functions will receive the table as their only argument.

9.2.4.1 reader

`LUATEX` will run this function whenever it needs a new input line from the file.

```
function(<table> env)
    return <string> line
end
```

Your function should return either a string or `nil`. The value `nil` signals that the end of file has occurred, and will make `TEX` call the optional `close` function next.

9.2.4.2 close

`LUATEX` will run this optional function when it decides to close the file.

```
function(<table> env)
end
```

Your function should not return any value.

9.3 Data processing callbacks

9.3.1 process_jobname

This callback allows you to change the jobname given by `\jobname` in `TEX` and `tex.jobname` in Lua. It does not affect the internal job name or the name of the output or log files.

```
function(<string> jobname)
    return <string> adjusted_jobname
```



end

The only argument is the actual job name; you should not use `tex.jobname` inside this function or infinite recursion may occur. If you return `nil`, L^AT_EX will pretend your callback never happened. This callback does not replace any internal code.

9.4 Node list processing callbacks

The description of nodes and node lists is in chapter 8.

9.4.1 `contribute_filter`

This callback is called when L^AT_EX adds contents to list:

```
function(<string> extrainfo)
end
```

The string reports the group code. From this you can deduce from what list you can give a treat.

VALUE	EXPLANATION
<code>pre_box</code>	interline material is being added
<code>pre_adjust</code>	<code>\vadjust</code> material is being added
<code>box</code>	a typeset box is being added (always called)
<code>adjust</code>	<code>\vadjust</code> material is being added

9.4.2 `buildpage_filter`

This callback is called whenever L^AT_EX is ready to move stuff to the main vertical list. You can use this callback to do specialized manipulation of the page building stage like imposition or column balancing.

```
function(<string> extrainfo)
end
```

The string `extrainfo` gives some additional information about what T_EX's state is with respect to the 'current page'. The possible values for the `buildpage_filter` callback are:

VALUE	EXPLANATION
<code>alignment</code>	a (partial) alignment is being added
<code>after_output</code>	an output routine has just finished
<code>new_graf</code>	the beginning of a new paragraph
<code>vmode_par</code>	<code>\par</code> was found in vertical mode
<code>hmode_par</code>	<code>\par</code> was found in horizontal mode
<code>insert</code>	an insert is added
<code>penalty</code>	a penalty (in vertical mode)
<code>before_display</code>	immediately before a display starts



after_display	a display is finished
end	LUAT _E X is terminating (it's all over)

9.4.3 build_page_insert

This callback is called when the pagebuilder adds an insert. There is not much control over this mechanism but this callback permits some last minute manipulations of the spacing before an insert, something that might be handy when for instance multiple inserts (types) are appended in a row.

```
function(<number> n, <number> i)
    return <number> register
end
```

with

VALUE	EXPLANATION
n	the insert class
i	the order of the insert

The return value is a number indicating the skip register to use for the prepended spacing. This permits for instance a different top space (when *i* equals one) and intermediate space (when *i* is larger than one). Of course you can mess with the insert box but you need to make sure that LUAT_EX is happy afterwards.

9.4.4 pre_linebreak_filter

This callback is called just before LUAT_EX starts converting a list of nodes into a stack of \hboxes, after the addition of \parfillskip.

```
function(<node> head, <string> groupcode)
    return true | false | <node> newhead
end
```

The string called groupcode identifies the nodelist's context within T_EX's processing. The range of possibilities is given in the table below, but not all of those can actually appear in pre_linebreak_filter, some are for the hpack_filter and vpack_filter callbacks that will be explained in the next two paragraphs.

VALUE	EXPLANATION
<empty>	main vertical list
hbox	\hbox in horizontal mode
adjusted_hbox	\hbox in vertical mode
vbox	\vbox
vtop	\vtop
align	\halign or \valign
disc	discretionaries



<code>insert</code>	packaging an insert
<code>vcenter</code>	<code>\vcenter</code>
<code>local_box</code>	<code>\localleftbox</code> or <code>\localrightbox</code>
<code>split_off</code>	top of a <code>\vsplit</code>
<code>split_keep</code>	remainder of a <code>\vsplit</code>
<code>align_set</code>	alignment cell
<code>fin_row</code>	alignment row

As for all the callbacks that deal with nodes, the return value can be one of three things:

- ▶ boolean true signals successful processing
- ▶ `<node>` signals that the ‘head’ node should be replaced by the returned node
- ▶ boolean false signals that the ‘head’ node list should be ignored and flushed from memory

This callback does not replace any internal code.

9.4.5 `linebreak_filter`

This callback replaces L^AT_EX’s line breaking algorithm.

```
function(<node> head, <boolean> is_display)
    return <node> newhead
end
```

The returned node is the head of the list that will be added to the main vertical list, the boolean argument is true if this paragraph is interrupted by a following math display.

If you return something that is not a `<node>`, L^AT_EX will apply the internal linebreak algorithm on the list that starts at `<head>`. Otherwise, the `<node>` you return is supposed to be the head of a list of nodes that are all allowed in vertical mode, and at least one of those has to represent a `hbox`. Failure to do so will result in a fatal error.

Setting this callback to false is possible, but dangerous, because it is possible you will end up in an unfixable ‘deadcycles loop’.

9.4.6 `append_to_vlist_filter`

This callback is called whenever L^AT_EX adds a box to a vertical list:

```
function(<node> box, <string> locationcode, <number prevdepth>,
    <boolean> mirrored)
    return list, prevdepth
end
```

It is ok to return nothing in which case you also need to flush the box or deal with it yourself. The `prevdepth` is also optional. Locations are `box`, `alignment`, `equation`, `equation_number` and `post_linebreak`.



9.4.7 post_linebreak_filter

This callback is called just after L^AT_EX has converted a list of nodes into a stack of \hboxes.

```
function(<node> head, <string> groupcode)
    return true | false | <node> newhead
end
```

This callback does not replace any internal code.

9.4.8 hpack_filter

This callback is called when T_EX is ready to start boxing some horizontal mode material. Math items and line boxes are ignored at the moment.

```
function(<node> head, <string> groupcode, <number> size,
        <string> packtype [, <string> direction] [, <node> attributelist])
    return true | false | <node> newhead
end
```

The packtype is either additional or exactly. If additional, then the size is a \hbox spread ... argument. If exactly, then the size is a \hbox to In both cases, the number is in scaled points.

The direction is either one of the three-letter direction specifier strings, or nil.

This callback does not replace any internal code.

9.4.9 vpack_filter

This callback is called when T_EX is ready to start boxing some vertical mode material. Math displays are ignored at the moment.

This function is very similar to the hpack_filter. Besides the fact that it is called at different moments, there is an extra variable that matches T_EX's \maxdepth setting.

```
function(<node> head, <string> groupcode, <number> size, <string> packtype,
        <number> maxdepth [, <string> direction] [, <node> attributelist]))
    return true | false | <node> newhead
end
```

This callback does not replace any internal code.

9.4.10 hpack_quality

This callback can be used to intercept the overfull messages that can result from packing a horizontal list (as happens in the par builder). The function takes a few arguments:

```
function(<string> incident, <number> detail, <node> head, <number> first,
        <number> last)
```



```

    return <node> whatever
end

```

The incident is one of `overflow`, `underfull`, `loose` or `tight`. The detail is either the amount of overflow in case of `overflow`, or the badness otherwise. The head is the list that is constructed (when protrusion or expansion is enabled, this is an intermediate list). Optionally you can return a node, for instance an `overflow` rule indicator. That node will be appended to the list (just like `TEX`'s own rule would).

9.4.11 `vpack_quality`

This callback can be used to intercept the `overflow` messages that can result from packing a vertical list (as happens in the page builder). The function takes a few arguments:

```

function(<string> incident, <number> detail, <node> head, <number> first,
        <number> last)
end

```

The incident is one of `overflow`, `underfull`, `loose` or `tight`. The detail is either the amount of overflow in case of `overflow`, or the badness otherwise. The head is the list that is constructed.

9.4.12 `process_rule`

This is an experimental callback. It can be used with rules of subtype 4 (user). The callback gets three arguments: the node, the width and the height. The callback can use `pdf.print` to write code to the PDF file but beware of not messing up the final result. No checking is done.

9.4.13 `pre_output_filter`

This callback is called when `TEX` is ready to start boxing the box 255 for `\output`.

```

function(<node> head, <string> groupcode, <number> size, <string> packtype,
        <number> maxdepth [, <string> direction])
    return true | false | <node> newhead
end

```

This callback does not replace any internal code.

9.4.14 `hyphenate`

```

function(<node> head, <node> tail)
end

```

No return values. This callback has to insert discretionary nodes in the node list it receives.

Setting this callback to `false` will prevent the internal discretionary insertion pass.



9.4.15 ligaturing

```
function(<node> head, <node> tail)
end
```

No return values. This callback has to apply ligaturing to the node list it receives.

You don't have to worry about return values because the head node that is passed on to the callback is guaranteed not to be a `glyph_node` (if need be, a temporary node will be prepended), and therefore it cannot be affected by the mutations that take place. After the callback, the internal value of the 'tail of the list' will be recalculated.

The next of head is guaranteed to be non-nil.

The next of tail is guaranteed to be nil, and therefore the second callback argument can often be ignored. It is provided for orthogonality, and because it can sometimes be handy when special processing has to take place.

Setting this callback to false will prevent the internal ligature creation pass.

You must not ruin the node list. For instance, the head normally is a local par node, and the tail a glue. Messing too much can push L^AT_EX into panic mode.

9.4.16 kerning

```
function(<node> head, <node> tail)
end
```

No return values. This callback has to apply kerning between the nodes in the node list it receives. See `ligaturing` for calling conventions.

Setting this callback to false will prevent the internal kern insertion pass.

You must not ruin the node list. For instance, the head normally is a local par node, and the tail a glue. Messing too much can push L^AT_EX into panic mode.

9.4.17 insert_local_par

Each paragraph starts with a local par node that keeps track of for instance the direction. You can hook a callback into the creator:

```
function(<node> local_par, <string> location)
end
```

There is no return value and you should make sure that the node stays valid as otherwise T_EX can get confused.

9.4.18 mlist_to_hlist

This callback replaces L^AT_EX's math list to node list conversion algorithm.

```
function(<node> head, <string> display_type, <boolean> need_penalties)
```



```
    return <node> newhead
end
```

The returned node is the head of the list that will be added to the vertical or horizontal list, the string argument is either 'text' or 'display' depending on the current math mode, the boolean argument is true if penalties have to be inserted in this list, false otherwise.

Setting this callback to false is bad, it will almost certainly result in an endless loop.

9.5 Information reporting callbacks

9.5.1 pre_dump

```
function()
end
```

This function is called just before dumping to a format file starts. It does not replace any code and there are neither arguments nor return values.

9.5.2 start_run

```
function()
end
```

This callback replaces the code that prints L^AT_EX's banner. Note that for successful use, this callback has to be set in the LUA initialization script, otherwise it will be seen only after the run has already started.

9.5.3 stop_run

```
function()
end
```

This callback replaces the code that prints L^AT_EX's statistics and 'output written to' messages. The engine can still do housekeeping and therefore you should not rely on this hook for postprocessing the PDF or log file.

9.5.4 show_error_hook

```
function()
end
```

This callback is run from inside the T_EX error function, and the idea is to allow you to do some extra reporting on top of what T_EX already does (none of the normal actions are removed). You may find some of the values in the status table useful. This callback does not replace any internal code.



9.5.5 show_error_message

```
function()  
end
```

This callback replaces the code that prints the error message. The usual interaction after the message is not affected.

9.5.6 show_lua_error_hook

```
function()  
end
```

This callback replaces the code that prints the extra LUA error message.

9.5.7 start_file

```
function(category,filename)  
end
```

This callback replaces the code that prints L^AT_EX's when a file is opened like (filename for regular files. The category is a number:

VALUE	MEANING
1	a normal data file, like a source
2	a font map coupling font names to resources
3	an image file (png, pdf, etc)
4	an embedded font subset
5	a fully embedded font

9.5.8 stop_file

```
function(category)  
end
```

This callback replaces the code that prints L^AT_EX's when a file is closed like the) for regular files.

9.5.9 wrapup_run

This callback is called after the PDF and log files are closed. Use it at your own risk.

9.6 Font-related callbacks

9.6.1 define_font

```
function(<string> name, <number> size)
```



```
    return <number> id
end
```

The string name is the filename part of the font specification, as given by the user.

The number size is a bit special:

- ▶ If it is positive, it specifies an ‘at size’ in scaled points.
- ▶ If it is negative, its absolute value represents a ‘scaled’ setting relative to the design size of the font.

The font can be defined with `font.define` which returns a font identifier that can be returned in the callback. So, contrary to L^AT_EX, in L^AMETAT_EX we only accept a number.

The internal structure of the font table that is passed to `font.define` is explained in chapter 6. That table is saved internally, so you can put extra fields in the table for your later LUA code to use. In alternative, `retval` can be a previously defined fontid. This is useful if a previous definition can be reused instead of creating a whole new font structure.

Setting this callback to `false` is pointless as it will prevent font loading completely but will nevertheless generate errors.



10 The T_EX related libraries

10.1 The lua library

10.1.1 Version information

This library contains two read-only items:

```
<string> v = lua.version  
<string> s = lua.startupfile
```

This returns the LUA version identifier string. The value is currently Lua 5.4.

10.1.2 Table allocators

Sometimes performance (and memory usage) can benefit a little from it preallocating a table with `newtable`:

```
<table> t = lua.newtable(100,5000)
```

This preallocates 100 hash entries and 5000 index entries. The `newindex` function create an indexed table with preset values:

```
<table> t = lua.newindex(2500,true)
```

10.1.3 Bytecode registers

LUA registers can be used to store LUA code chunks. The accepted values for assignments are functions and `nil`. Likewise, the retrieved value is either a function or `nil`.

```
lua.bytecode[<number> n] = <function> f  
lua.bytecode[<number> n]()
```

The contents of the `lua.bytecode` array is stored inside the format file as actual LUA bytecode, so it can also be used to preload LUA code. The function must not contain any upvalues. The associated function calls are:

```
<function> f = lua.getbytecode(<number> n)  
lua.setbytecode(<number> n, <function> f)
```

Note: Since a LUA file loaded using `loadfile(filename)` is essentially an anonymous function, a complete file can be stored in a bytecode register like this:

```
lua.bytecode[n] = loadfile(filename)
```

Now all definitions (functions, variables) contained in the file can be created by executing this bytecode register:



```
lua.bytecode[n]()
```

Note that the path of the file is stored in the LUA bytecode to be used in stack backtraces and therefore dumped into the format file if the above code is used in `INITTEX`. If it contains private information, i.e. the user name, this information is then contained in the format file as well. This should be kept in mind when preloading files into a bytecode register in `INITTEX`.

10.1.4 Chunk name registers

There is an array of 65536 (0-65535) potential chunk names for use with the `\directlua` and `\latelua` primitives.

```
lua.name[<number> n] = <string> s
<string> s = lua.name[<number> n]
```

If you want to unset a LUA name, you can assign `nil` to it. The function accessors are:

```
lua.setluaname(<string> s,<number> n)
<string> s = lua.getluaname(<number> n)
```

10.1.5 Introspection

The `getstacktop` function return a number indicating how full the LUA stack is. This function only makes sense as breakpoint when checking some mechanism going haywire.

There are four time related helpers. The `getruntime` function returns the time passed since startup. The `getcurrenttime` does what its name says. Just play with them to see how it pays off. The `getpreciseticks` returns a number that can be used later, after a similar call, to get a difference. The `getpreciseseconds` function gets such a tick (delta) as argument and returns the number of seconds. Ticks can differ per operating system, but one always creates a reference first and then deltas to this reference.

10.2 The status library

This contains a number of run-time configuration items that you may find useful in message reporting, as well as an iterator function that gets all of the names and values as a table.

```
<table> info = status.list()
```

The keys in the table are the known items, the value is the current value. Almost all of the values in `status` are fetched through a metatable at run-time whenever they are accessed, so you cannot use `pairs` on `status`, but you *can* use `pairs` on `info`, of course. If you do not need the full list, you can also ask for a single item by using its name as an index into `status`. The current list is:

KEY	EXPLANATION
<code>banner</code>	terminal display banner
<code>best_page_break</code>	the current best break (a node)



buf_size	current allocated size of the line buffer
callbacks	total number of executed callbacks so far
cs_count	number of control sequences
dest_names_size	PDF destination table size
dvi_gone	written DVI bytes
dvi_ptr	not yet written DVI bytes
dyn_used	token (multi-word) memory in use
filename	name of the current input file
fix_mem_end	maximum number of used tokens
fix_mem_min	minimum number of allocated words for tokens
fix_mem_max	maximum number of allocated words for tokens
font_ptr	number of active fonts
hash_extra	extra allowed hash
hash_size	size of hash
indirect_callbacks	number of those that were themselves a result of other callbacks (e.g. file readers)
ini_version	true if this is an INIT _{TeX} run
init_pool_ptr	INIT _{TeX} string pool index
init_str_ptr	number of INIT _{TeX} strings
input_ptr	the level of input we're at
inputid	numeric id of the current input
largest_used_mark	max referenced marks class
lasterrorcontext	last error context string (with newlines)
lasterrorstring	last _{TeX} error string
lastluaerrorstring	last LUA error string
lastwarningstring	last warning tag, normally an indication of in what part
lastwarningtag	last warning string
linenumber	location in the current input file
log_name	name of the log file
luabytecode_bytes	number of bytes in LUA bytecode registers
luabytecodes	number of active LUA bytecode registers
luastate_bytes	number of bytes in use by LUA interpreters
luatex_engine	the LUAT _{TeX} engine identifier
luatex_hashchars	length to which LUA hashes strings (2^n)
luatex_hashtype	the hash method used (in LUAJIT _{TeX})
luatex_version	the LUAT _{TeX} version number
luatex_revision	the LUAT _{TeX} revision string
max_buf_stack	max used buffer position
max_in_stack	max used input stack entries
max_nest_stack	max used nesting stack entries
max_param_stack	max used parameter stack entries
max_save_stack	max used save stack entries
max_strings	maximum allowed strings
nest_size	nesting stack size
node_mem_usage	a string giving insight into currently used nodes
obj_ptr	max PDF object pointer



<code>obj_tab_size</code>	PDF object table size
<code>output_active</code>	true if the <code>\output</code> routine is active
<code>output_file_name</code>	name of the PDF or DVI file
<code>param_size</code>	parameter stack size
<code>pdf_dest_names_ptr</code>	max PDF destination pointer
<code>pdf_gone</code>	written PDF bytes
<code>pdf_mem_ptr</code>	max PDF memory used
<code>pdf_mem_size</code>	PDF memory size
<code>pdf_os_cntr</code>	max PDF object stream pointer
<code>pdf_os_objidx</code>	PDF object stream index
<code>pdf_ptr</code>	not yet written PDF bytes
<code>pool_ptr</code>	string pool index
<code>pool_size</code>	current size allocated for string characters
<code>save_size</code>	save stack size
<code>shell_escape</code>	0 means disabled, 1 means anything is permitted, and 2 is restricted
<code>safer_option</code>	1 means safer is enforced
<code>kpse_used</code>	1 means that kpse is used
<code>stack_size</code>	input stack size
<code>str_ptr</code>	number of strings
<code>total_pages</code>	number of written pages
<code>var_mem_max</code>	number of allocated words for nodes
<code>var_used</code>	variable (one-word) memory in use
<code>lc_collate</code>	the value of <code>LC_COLLATE</code> at startup time (becomes C at startup)
<code>lc_ctype</code>	the value of <code>LC_CTYPE</code> at startup time (becomes C at startup)
<code>lc_numeric</code>	the value of <code>LC_NUMERIC</code> at startup time

The error and warning messages can be wiped with the `resetmessages` function. A return value can be set with `setexitcode`.

10.3 The tex library

10.3.1 Introduction

The `tex` table contains a large list of virtual internal \TeX parameters that are partially writable. The designation ‘virtual’ means that these items are not properly defined in `LUA`, but are only frontends that are handled by a metatable that operates on the actual \TeX values. As a result, most of the `LUA` table operators (like `pairs` and `#`) do not work on such items.

At the moment, it is possible to access almost every parameter that you can use after `\the`, is a single tokens or is sort of special in \TeX . This excludes parameters that need extra arguments, like `\the\scriptfont`. The subset comprising simple integer and dimension registers are writable as well as readable (like `\tracingcommands` and `\parindent`).

10.3.2 Internal parameter values, set and get

For all the parameters in this section, it is possible to access them directly using their names as index in the `tex` table, or by using one of the functions `tex.get` and `tex.set`.



The exact parameters and return values differ depending on the actual parameter, and so does whether `tex.set` has any effect. For the parameters that *can* be set, it is possible to use `global` as the first argument to `tex.set`; this makes the assignment global instead of local.

```
tex.set ([ "global", ] <string> n, ...)
... = tex.get (<string> n)
```

Glue is kind of special because there are five values involved. The return value is a `glue_spec` node but when you pass `false` as last argument to `tex.get` you get the width of the glue and when you pass `true` you get all five values. Otherwise you get a node which is a copy of the internal value so you are responsible for its freeing at the LUA end. When you set a glue quantity you can either pass a `glue_spec` or upto five numbers. If you pass `true` to get you get 5 values returned for a glue and when you pass `false` you only get the width returned.

10.3.2.1 Integer parameters

The integer parameters accept and return LUA numbers. These are read-write:

<code>tex.adjdemerits</code>	<code>tex.localinterlinepenalty</code>
<code>tex.binoppenalty</code>	<code>tex.looseness</code>
<code>tex.brokenpenalty</code>	<code>tex.mag</code>
<code>tex.catcodetable</code>	<code>tex.maxdeadcycles</code>
<code>tex.clubpenalty</code>	<code>tex.month</code>
<code>tex.day</code>	<code>tex.newlinechar</code>
<code>tex.defaultthyphenchar</code>	<code>tex.outputpenalty</code>
<code>tex.defaultskewchar</code>	<code>tex.pausing</code>
<code>tex.delimiterfactor</code>	<code>tex.postdisplaypenalty</code>
<code>tex.displaywidowpenalty</code>	<code>tex.predisplaydirection</code>
<code>tex.doublehyphendemerits</code>	<code>tex.predisplaypenalty</code>
<code>tex.endlinechar</code>	<code>tex.pretolerance</code>
<code>tex.errorcontextlines</code>	<code>tex.relpenny</code>
<code>tex.escapechar</code>	<code>tex.righthyphenmin</code>
<code>tex.exhyphenpenalty</code>	<code>tex.savinghyphcodes</code>
<code>tex.fam</code>	<code>tex.savingvdiscards</code>
<code>tex.finalhyphendemerits</code>	<code>tex.showboxbreadth</code>
<code>tex.floatingpenalty</code>	<code>tex.showboxdepth</code>
<code>tex.globaldefs</code>	<code>tex.time</code>
<code>tex.hangafter</code>	<code>tex.tolerance</code>
<code>tex.hbadness</code>	<code>tex.tracingassigns</code>
<code>tex.holdinginserts</code>	<code>tex.tracingcommands</code>
<code>tex.hyphenpenalty</code>	<code>tex.tracinggroups</code>
<code>tex.interlinepenalty</code>	<code>tex.tracingifs</code>
<code>tex.language</code>	<code>tex.tracinglostchars</code>
<code>tex.lastlinefit</code>	<code>tex.tracingmacros</code>
<code>tex.lefthyphenmin</code>	<code>tex.tracingnesting</code>
<code>tex.linepenalty</code>	<code>tex.tracingonline</code>
<code>tex.localbrokenpenalty</code>	<code>tex.tracingoutput</code>



<code>tex.tracingpages</code>	<code>tex.uchyph</code>
<code>tex.tracingparagraphs</code>	<code>tex.vbadness</code>
<code>tex.tracingrestores</code>	<code>tex.widowpenalty</code>
<code>tex.tracingscantokens</code>	<code>tex.year</code>
<code>tex.tracingstats</code>	

These are read-only:

<code>tex.deadcycles</code>	<code>tex.parshape</code>	<code>tex.spacefactor</code>
<code>tex.insertpenalties</code>	<code>tex.prevgraf</code>	

10.3.2.2 Dimension parameters

The dimension parameters accept LUA numbers (signifying scaled points) or strings (with included dimension). The result is always a number in scaled points. These are read-write:

<code>tex.boxmaxdepth</code>	<code>tex.mathsurround</code>	<code>tex.parindent</code>
<code>tex.delimitershortfall</code>	<code>tex.maxdepth</code>	<code>tex.predisplaysize</code>
<code>tex.displayindent</code>	<code>tex.nulldelimiterspace</code>	<code>tex.scriptsace</code>
<code>tex.displaywidth</code>	<code>tex.overfullrule</code>	<code>tex.splitmaxdepth</code>
<code>tex.emergencystretch</code>	<code>tex.pagebottomoffset</code>	<code>tex.vfuzz</code>
<code>tex.hangindent</code>	<code>tex.pageheight</code>	<code>tex.voffset</code>
<code>tex.hfuzz</code>	<code>tex.pageleftoffset</code>	<code>tex.vsize</code>
<code>tex.hoffset</code>	<code>tex.pagerightoffset</code>	<code>tex.prevdepth</code>
<code>tex.hsize</code>	<code>tex.pagetopoffset</code>	<code>tex.prevgraf</code>
<code>tex.lineskiplimit</code>	<code>tex.pagewidth</code>	<code>tex.spacefactor</code>

These are read-only:

<code>tex.pagedepth</code>	<code>tex.pagefilstretch</code>	<code>tex.pagestretch</code>
<code>tex.pagefilllstretch</code>	<code>tex.pagegoal</code>	<code>tex.pagetotal</code>
<code>tex.pagefillstretch</code>	<code>tex.pageshrink</code>	

Beware: as with all LUA tables you can add values to them. So, the following is valid:

```
tex.foo = 123
```

When you access a T_EX parameter a look up takes place. For read-only variables that means that you will get something back, but when you set them you create a new entry in the table thereby making the original invisible.

There are a few special cases that we make an exception for: `prevdepth`, `prevgraf` and `spacefactor`. These normally are accessed via the `tex.nest` table:

```
tex.nest[tex.nest.ptr].prevdepth = p
tex.nest[tex.nest.ptr].spacefactor = s
```

However, the following also works:



```
tex.prevdepth = p
tex.spacefactor = s
```

Keep in mind that when you mess with node lists directly at the LUA end you might need to update the top of the nesting stack's prevdepth explicitly as there is no way L^AT_EX can guess your intentions. By using the accessor in the tex tables, you get and set the values at the top of the nesting stack.

10.3.2.3 Direction parameters

The direction states can be queried and set with:

tex.gettextdir()	tex.getpardir()	tex.setmathdir(<number>)
tex.getlinedir()	tex.settextdir(<number>)	tex.setpardir(<number>)
tex.getmathdir()	tex.setlinedir(<number>)	

and also with:

tex.textdirection	tex.mathdirection
tex.linedirection	tex.pardirection

10.3.2.4 Glue parameters

The glue parameters accept and return a userdata object that represents a glue_spec node.

tex.abovedisplayskip	tex.leftskip	tex.spaceskip
tex.abovedisplayskip	tex.lineskip	tex.splittopskip
tex.baselineskip	tex.parfillskip	tex.tabskip
tex.belowdisplayskip	tex.parskip	tex.topskip
tex.belowdisplayskip	tex.rightskip	tex.xspaceskip

10.3.2.5 Muglue parameters

All muglue parameters are to be used read-only and return a LUA string.

tex.medmuskip	tex.thickmuskip	tex.thinmuskip
---------------	-----------------	----------------

10.3.2.6 Tokenlist parameters

The tokenlist parameters accept and return LUA strings. LUA strings are converted to and from token lists using \the \toks style expansion: all category codes are either space (10) or other (12). It follows that assigning to some of these, like 'tex.output', is actually useless, but it feels bad to make exceptions in view of a coming extension that will accept full-blown token strings.

tex.errhelp	tex.everyeof	tex.everymath
tex.everycr	tex.everyhbox	tex.everypar
tex.everydisplay	tex.everyjob	tex.everyvbox



tex.output

10.3.3 Convert commands

All ‘convert’ commands are read-only and return a LUA string. The supported commands at this moment are:

tex.eTeXVersion	tex.fontname(number)
tex.eTeXrevision	tex.uniformdeviate(number)
tex.formatname	tex.number(number)
tex.jobname	tex.romannumeral(number)
tex.luatexbanner	tex.fontidentifier(number)
tex.luatexrevision	

If you are wondering why this list looks haphazard; these are all the cases of the ‘convert’ internal command that do not require an argument, as well as the ones that require only a simple numeric value. The special (LUA-only) case of tex.fontidentifier returns the csname string that matches a font id number (if there is one).

10.3.4 Last item commands

All ‘last item’ commands are read-only and return a number. The supported commands at this moment are:

tex.lastpenalty	tex.lasttypos	tex.currentgrouptype
tex.lastkern	tex.randomseed	tex.currentiflevel
tex.lastskip	tex.luatexversion	tex.currentifttype
tex.lastnodetype	tex.eTeXminorversion	tex.currentifbranch
tex.inputlineno	tex.eTeXversion	
tex.lastxpos	tex.currentgrouplevel	

10.3.5 Accessing registers: set*, get* and is*

T_EX’s attributes (\attribute), counters (\count), dimensions (\dimen), skips (\skip, \muskip) and token (\toks) registers can be accessed and written to using two times five virtual sub-tables of the tex table:

tex.attribute	tex.skip	tex.muglue
tex.count	tex.glue	tex.toks
tex.dimen	tex.muskip	

It is possible to use the names of relevant \attributedef, \countdef, \dimendef, \skipdef, or \toksdef control sequences as indices to these tables:

```
tex.count.scratchcounter = 0
```



```
enormous = tex.dimen['maxdimen']
```

In this case, L^AT_EX looks up the value for you on the fly. You have to use a valid `\countdef` (or `\attributedef`, or `\dimendef`, or `\skipdef`, or `\toksdef`), anything else will generate an error (the intent is to eventually also allow `<chardef tokens>` and even macros that expand into a number).

- ▶ The count registers accept and return LUA numbers.
- ▶ The dimension registers accept LUA numbers (in scaled points) or strings (with an included absolute dimension; `em` and `ex` and `px` are forbidden). The result is always a number in scaled points.
- ▶ The token registers accept and return LUA strings. LUA strings are converted to and from token lists using `\the \toks` style expansion: all category codes are either space (10) or other (12).
- ▶ The skip registers accept and return `glue_spec` userdata node objects (see the description of the node interface elsewhere in this manual).
- ▶ The glue registers are just skip registers but instead of userdata are verbose.
- ▶ Like the counts, the attribute registers accept and return LUA numbers.

As an alternative to array addressing, there are also accessor functions defined for all cases, for example, here is the set of possibilities for `\skip` registers:

```
tex.setskip ([ "global", ] <number> n, <node> s )
tex.setskip ([ "global", ] <string> s, <node> s )
<node> s = tex.getskip (<number> n)
<node> s = tex.getskip (<string> s)
```

We have similar setters for `count`, `dimen`, `muskip`, and `toks`. Counters and `dimen` are represented by numbers, `skips` and `muskip`s by nodes, and `toks` by strings.

Again the glue variants are not using the `glue-spec` userdata nodes. The `setglue` function accepts upto 5 arguments: `width`, `stretch`, `shrink`, `stretch_order` and `shrink_order` and the `getglue` function reports them, unless the second argument is `false` in which case only the `width` is returned.

Here is an example using a threesome:

```
local d = tex.getdimen("foo")
if tex.isdimen("bar") then
    tex.setdimen("bar",d)
end
```

There are six extra `skip` (glue) related helpers:

```
tex.setglue ([ "global", ] <number> n,
    width, stretch, shrink, stretch_order, shrink_order)
tex.setglue ([ "global", ] <string> s,
    width, stretch, shrink, stretch_order, shrink_order)
width, stretch, shrink, stretch_order, shrink_order =
    tex.getglue (<number> n)
```



```
width, stretch, shrink, stretch_order, shrink_order =
    tex.getglue (<string> s)
```

The other two are `tex.setmuglue` and `tex.getmuglue`.

There are such helpers for `dimen`, `count`, `skip`, `muskip`, `box` and `attribute` registers but the glue ones are special because they have to deal with more properties.

As with the general `get` and `set` function discussed before, for the `skip` registers `getskip` returns a node and `getglue` returns numbers, while `setskip` accepts a node and `setglue` expects upto 5 numbers. Again, when you pass `false` as second argument to `getglue` you only get the width returned. The same is true for the `mu` variants `getmuskip`, `setmuskip`, `getmuskip` and `setmuskip`.

For tokens registers we have an alternative where a catcode table is specified:

```
tex.scantoks(0,3,"$e=mc^2$")
tex.scantoks("global",0,3,"$\int\limits^1_2$")
```

In the function-based interface, it is possible to define values globally by using the string `global` as the first function argument.

10.3.6 Character code registers: `[get|set]*code[s]`

\TeX 's character code tables (`\lccode`, `\uccode`, `\sfcode`, `\catcode`, `\mathcode`, `\delcode`) can be accessed and written to using six virtual subtables of the `tex` table

<code>tex.lccode</code>	<code>tex.sfcode</code>	<code>tex.mathcode</code>
<code>tex.uccode</code>	<code>tex.catcode</code>	<code>tex.delcode</code>

The function call interfaces are roughly as above, but there are a few twists. `sfcodes` are the simple ones:

```
tex.setsfcode (["global",] <number> n, <number> s)
<number> s = tex.getsfcode (<number> n)
```

The function call interface for `lccode` and `uccode` additionally allows you to set the associated sibling at the same time:

```
tex.setlccode (["global"], <number> n, <number> lc)
tex.setlccode (["global"], <number> n, <number> lc, <number> uc)
<number> lc = tex.getlccode (<number> n)
tex.setuccode (["global"], <number> n, <number> uc)
tex.setuccode (["global"], <number> n, <number> uc, <number> lc)
<number> uc = tex.getuccode (<number> n)
```

The function call interface for `catcode` also allows you to specify a category table to use on assignment or on query (default in both cases is the current one):

```
tex.setcatcode (["global"], <number> n, <number> c)
tex.setcatcode (["global"], <number> cattable, <number> n, <number> c)
```




```

<number> lc = tex.getcatcode (<number> n)
<number> lc = tex.getcatcode (<number> cattable, <number> n)

```

The interfaces for `delcode` and `mathcode` use small array tables to set and retrieve values:

```

tex.setmathcode (["global"], <number> n, <table> mval )
<table> mval = tex.getmathcode (<number> n)
tex.setdelcode (["global"], <number> n, <table> dval )
<table> dval = tex.getdelcode (<number> n)

```

Where the table for `mathcode` is an array of 3 numbers, like this:

```

{
    <number> class,
    <number> family,
    <number> character
}

```

And the table for `delcode` is an array with 4 numbers, like this:

```

{
    <number> small_fam,
    <number> small_char,
    <number> large_fam,
    <number> large_char
}

```

You can also avoid the table:

```

tex.setmathcode (["global"], <number> n, <number> class,
    <number> family, <number> character)
class, family, char =
    tex.getmathcodes (<number> n)
tex.setdelcode (["global"], <number> n, <number> smallfam,
    <number> smallchar, <number> largefam, <number> largechar)
smallfam, smallchar, largefam, largechar =
    tex.getdelcodes (<number> n)

```

Normally, the third and fourth values in a delimiter code assignment will be zero according to `\Udelcode` usage, but the returned table can have values there (if the delimiter code was set using `\delcode`, for example). Unset `delcode`'s can be recognized because `dval[1]` is `-1`.

10.3.7 Box registers: `[get|set]box`

It is possible to set and query actual boxes, coming for instance from `\hbox`, `\vbox` or `\vtop`, using the node interface as defined in the node library:

```
tex.box
```

for array access, or



```
tex.setbox(["global",] <number> n, <node> s)
tex.setbox(["global",] <string> cs, <node> s)
<node> n = tex.getbox(<number> n)
<node> n = tex.getbox(<string> cs)
```

for function-based access. In the function-based interface, it is possible to define values globally by using the string `global` as the first function argument.

Be warned that an assignment like

```
tex.box[0] = tex.box[2]
```

does not copy the node list, it just duplicates a node pointer. If `\box2` will be cleared by \TeX commands later on, the contents of `\box0` becomes invalid as well. To prevent this from happening, always use `node.copy_list` unless you are assigning to a temporary variable:

```
tex.box[0] = node.copy_list(tex.box[2])
```

10.3.8 triggerbuildpage

You should not expect too much from the `triggerbuildpage` helpers because often \TeX doesn't do much if it thinks nothing has to be done, but it might be useful for some applications. It just does as it says it calls the internal function that build a page, given that there is something to build.

10.3.9 splitbox

You can split a box:

```
local vlist = tex.splitbox(n,height,mode)
```

The remainder is kept in the original box and a packaged vlist is returned. This operation is comparable to the `\vsplit` operation. The mode can be `additional` or `exactly` and concerns the split off box.

10.3.10 Accessing math parameters: [get|set]math

It is possible to set and query the internal math parameters using:

```
tex.setmath(["global",] <string> n, <string> t, <number> n)
<number> n = tex.getmath(<string> n, <string> t)
```

As before an optional first parameter `global` indicates a global assignment.

The first string is the parameter name minus the leading 'Umath', and the second string is the style name minus the trailing 'style'. Just to be complete, the values for the math parameter name are:

`quad`

`axis`

`operatorsize`



overbarkern	overbarrule	overbarvgap	
underbarkern	underbarrule	underbarvgap	
radicalkern	radicalrule	radicalvgap	
radicaldegreebefore	radicaldegreeafter	radicaldegreeraise	
stackvgap	stacknumup	stackdenomdown	
fractionrule	fractionnumvgap	fractionnumup	
fractiondenomvgap	fractiondenomdown	fractiondelsize	
limitabovevgap	limitabovebgap	limitabovekern	
limitbelowvgap	limitbelowbgap	limitbelowkern	
underdelimitervgap	underdelimiterbgap		
overdelimitervgap	overdelimiterbgap		
subshiftdrop	supshiftdrop	subshiftdown	
subsupshiftdown	subtopmax	supshiftup	
supbottommin	supsubbottommax	subsupvgap	
spaceafterscript	connectoroverlapmin		
ordordspacing	ordopspacing	ordbinspacing	ordrelspacing
ordopenspacing	ordclosespacing	ordpunctspacing	ordinnerspacing
opordspacing	opopspacing	opbinspacing	oprelspacing
opopenspacing	opclosespacing	oppunctspacing	opinnerspacing
binordspacing	binopspacing	binbinspacing	binrelspacing
binopenspacing	binclosespacing	binpunctspacing	bininnerspacing
relordspacing	relopspacing	relbinspacing	relrelspacing
relopenspacing	relclosespacing	relpunctspacing	relinnerspacing
openordspacing	openopspacing	openbinspacing	openrelspacing
openopenspacing	openclosespacing	openpunctspacing	openinnerspacing
closeordspacing	closeopspacing	closebinspacing	closerelspacing
closeopenspacing	closeclosespacing	closepunctspacing	closeinnerspacing
punctordspacing	punctopspacing	punctbinspacing	punctrelspacing
punctopenspacing	punctclosespacing	punctpunctspacing	punctinnerspacing
innerordspacing	inneropspacing	innerbinspacing	innerrelspacing
inneropenspacing	innerclosespacing	innerpunctspacing	innerinnerspacing

The values for the style parameter are:

display	crampeddisplay
text	crampedtext
script	crampedscript
scriptscript	crampedscriptscript

The value is either a number (representing a dimension or number) or a glue spec node representing a muskip for `ordordspacing` and similar spacing parameters.

10.3.11 Special list heads: `[get|set]list`

The virtual table `tex.lists` contains the set of internal registers that keep track of building page lists.



FIELD	EXPLANATION
page_ins_head	circular list of pending insertions
contrib_head	the recent contributions
page_head	the current page content
hold_head	used for held-over items for next page
adjust_head	head of the current \vadjust list
pre_adjust_head	head of the current \vadjust pre list
page_discards_head	head of the discarded items of a page break
split_discards_head	head of the discarded items in a vsplit

The getter and setter functions are `getlist` and `setlist`. You have to be careful with what you set as \TeX can have expectations with regards to how a list is constructed or in what state it is.

10.3.12 Semantic nest levels: `getnest` and `ptr`

The virtual table `nest` contains the currently active semantic nesting state. It has two main parts: a zero-based array of userdata for the semantic nest itself, and the numerical value `ptr`, which gives the highest available index. Neither the array items in `nest[]` nor `ptr` can be assigned to (as this would confuse the typesetting engine beyond repair), but you can assign to the individual values inside the array items, e.g. `tex.nest[tex.nest.ptr].prevdepth`.

`tex.nest[tex.nest.ptr]` is the current nest state, `nest[0]` the outermost (main vertical list) level. The getter function is `getnest`. You can pass a number (which gives you a list), nothing or `top`, which returns the topmost list, or the string `ptr` which gives you the index of the topmost list.

The known fields are:

KEY	TYPE	MODES	EXPLANATION
mode	number	all	the meaning of these numbers depends on the engine and sometimes even the version; you can use <code>tex.getmodevalues()</code> to get the mapping: positive values signal vertical, horizontal and math mode, while negative values indicate inner and inline variants
modeline	number	all	source input line where this mode was entered in, negative inside the output routine
head	node	all	the head of the current list
tail	node	all	the tail of the current list
prevgraf	number	vmode	number of lines in the previous paragraph
prevdepth	number	vmode	depth of the previous paragraph
spacefactor	number	hmode	the current space factor
dirs	node	hmode	used for temporary storage by the line break algorithm
noad	node	mmode	used for temporary storage of a pending fraction numerator, for <code>\over</code> etc.
delimptr	node	mmode	used for temporary storage of the previous math delimiter, for <code>\middle</code>
mathdir	boolean	mmode	true when during math processing the <code>\mathdir</code> is not the same as the surrounding <code>\texdir</code>



10.3.13 Print functions

The `tex` table also contains the three print functions that are the major interface from LUA scripting to \TeX . The arguments to these three functions are all stored in an in-memory virtual file that is fed to the \TeX scanner as the result of the expansion of `\directlua`.

The total amount of returnable text from a `\directlua` command is only limited by available system RAM. However, each separate printed string has to fit completely in \TeX 's input buffer. The result of using these functions from inside callbacks is undefined at the moment.

10.3.13.1 `print`

```
tex.print(<string> s, ...)
tex.print(<number> n, <string> s, ...)
tex.print(<table> t)
tex.print(<number> n, <table> t)
```

Each string argument is treated by \TeX as a separate input line. If there is a table argument instead of a list of strings, this has to be a consecutive array of strings to print (the first non-string value will stop the printing process).

The optional parameter can be used to print the strings using the catcode regime defined by `\catcodetable n`. If `n` is `-1`, the currently active catcode regime is used. If `n` is `-2`, the resulting catcodes are the result of `\the \toks`: all category codes are 12 (other) except for the space character, that has category code 10 (space). Otherwise, if `n` is not a valid catcode table, then it is ignored, and the currently active catcode regime is used instead.

The very last string of the very last `tex.print` command in a `\directlua` will not have the `\newlinechar` appended, all others do.

10.3.13.2 `sprint`

```
tex.sprint(<string> s, ...)
tex.sprint(<number> n, <string> s, ...)
tex.sprint(<table> t)
tex.sprint(<number> n, <table> t)
```

Each string argument is treated by \TeX as a special kind of input line that makes it suitable for use as a partial line input mechanism:

- ▶ \TeX does not switch to the 'new line' state, so that leading spaces are not ignored.
- ▶ No `\newlinechar` is inserted.
- ▶ Trailing spaces are not removed. Note that this does not prevent \TeX itself from eating spaces as result of interpreting the line. For example, in

```
before\directlua{tex.sprint("\relax")tex.sprint(" inbetween")}after
```



the space before `in` between will be gobbled as a result of the ‘normal’ scanning of `\relax`.

If there is a table argument instead of a list of strings, this has to be a consecutive array of strings to print (the first non-string value will stop the printing process).

The optional argument sets the catcode regime, as with `tex.print`. This influences the string arguments (or numbers turned into strings).

Although this needs to be used with care, you can also pass token or node userdata objects. These get injected into the stream. Tokens had best be valid tokens, while nodes need to be around when they get injected. Therefore it is important to realize the following:

- ▶ When you inject a token, you need to pass a valid token userdata object. This object will be collected by LUA when it no longer is referenced. When it gets printed to \TeX the token itself gets copied so there is no interference with the LUA garbage collection. You manage the object yourself. Because tokens are actually just numbers, there is no real extra overhead at the \TeX end.
- ▶ When you inject a node, you need to pass a valid node userdata object. The node related to the object will not be collected by LUA when it no longer is referenced. It lives on at the \TeX end in its own memory space. When it gets printed to \TeX the node reference is used assuming that node stays around. There is no LUA garbage collection involved. Again, you manage the object yourself. The node itself is freed when \TeX is done with it.

If you consider the last remark you might realize that we have a problem when a printed mix of strings, tokens and nodes is reused. Inside \TeX the sequence becomes a linked list of input buffers. So, `"123"` or `"\foo{123}"` gets read and parsed on the fly, while `<token userdata>` already is tokenized and effectively is a token list now. A `<node userdata>` is also tokenized into a token list but it has a reference to a real node. Normally this goes fine. But now assume that you store the whole lot in a macro: in that case the tokenized node can be flushed many times. But, after the first such flush the node is used and its memory freed. You can prevent this by using copies which is controlled by setting `\luacopyinputnodes` to a non-zero value. This is one of these fuzzy areas you have to live with if you really mess with these low level issues.

10.3.13.3 `tprint`

```
tex.tprint({<number> n, <string> s, ...}, {...})
```

This function is basically a shortcut for repeated calls to `tex.sprint(<number> n, <string> s, ...)`, once for each of the supplied argument tables.

10.3.13.4 `cprint`

This function takes a number indicating the to be used catcode, plus either a table of strings or an argument list of strings that will be pushed into the input stream.

```
tex.cprint( 1," 1: ${\\foo}") tex.print("\\par") -- a lot of \bgroup s
tex.cprint( 2," 2: ${\\foo}") tex.print("\\par") -- matching \egroup s
tex.cprint( 9," 9: ${\\foo}") tex.print("\\par") -- all get ignored
tex.cprint(10,"10: ${\\foo}") tex.print("\\par") -- all become spaces
```



```
tex.cprint(11,"11: ${\\foo}") tex.print("\\par") -- letters
tex.cprint(12,"12: ${\\foo}") tex.print("\\par") -- other characters
tex.cprint(14,"12: ${\\foo}") tex.print("\\par") -- comment triggers
```

10.3.13.5 write

```
tex.write(<string> s, ...)
tex.write(<table> t)
```

Each string argument is treated by T_EX as a special kind of input line that makes it suitable for use as a quick way to dump information:

- ▶ All catcodes on that line are either ‘space’ (for ‘ ’) or ‘character’ (for all others).
- ▶ There is no `\endlinechar` appended.

If there is a table argument instead of a list of strings, this has to be a consecutive array of strings to print (the first non-string value will stop the printing process).

10.3.14 Helper functions

10.3.14.1 round

```
<number> n = tex.round(<number> o)
```

Rounds LUA number `o`, and returns a number that is in the range of a valid T_EX register value. If the number starts out of range, it generates a ‘number too big’ error as well.

10.3.14.2 scale

```
<number> n = tex.scale(<number> o, <number> delta)
<table> n = tex.scale(table o, <number> delta)
```

Multiplies the LUA numbers `o` and `delta`, and returns a rounded number that is in the range of a valid T_EX register value. In the table version, it creates a copy of the table with all numeric top-level values scaled in that manner. If the multiplied number(s) are of range, it generates ‘number too big’ error(s) as well.

Note: the precision of the output of this function will depend on your computer’s architecture and operating system, so use with care! An interface to L^AT_EX’s internal, 100% portable scale function will be added at a later date.

10.3.14.3 number and romannumeral

These are the companions to the primitives `\number` and `\romannumeral`. They can be used like:

```
tex.print(tex.romannumeral(123))
```



10.3.14.4 fontidentifier and fontname

The first one returns the name only, the second one reports the size too.

```
tex.print(tex.fontname(tex.fontname))
tex.print(tex.fontname(tex.fontidentifier))
```

10.3.14.5 sp

```
<number> n = tex.sp(<number> o)
<number> n = tex.sp(<string> s)
```

Converts the number *o* or a string *s* that represents an explicit dimension into an integer number of scaled points.

For parsing the string, the same scanning and conversion rules are used that L^AT_EX would use if it was scanning a dimension specifier in its T_EX-like input language (this includes generating errors for bad values), expect for the following:

1. only explicit values are allowed, control sequences are not handled
2. infinite dimension units (*fil...*) are forbidden
3. *mu* units do not generate an error (but may not be useful either)

10.3.14.6 tex.getlinenumber and tex.setlinenumber

You can mess with the current line number:

```
local n = tex.getlinenumber()
tex.setlinenumber(n+10)
```

which can be shortcut to:

```
tex.setlinenumber(10,true)
```

This might be handy when you have a callback that read numbers from a file and combines them in one line (in which case an error message probably has to refer to the original line). Interference with T_EX's internal handling of numbers is of course possible.

10.3.14.7 error, show_context and gethelptext

```
tex.error(<string> s)
tex.error(<string> s, <table> help)
<string> s = tex.gethelptext()
```

This creates an error somewhat like the combination of `\errhelp` and `\errmessage` would. During this error, deletions are disabled.

The array part of the help table has to contain strings, one for each line of error help.

In case of an error the `show_context` function will show the current context where we're at (in the expansion).



10.3.14.8 getfamilyoffont

When you pass a proper family identifier the next helper will return the font currently associated with it.

```
<integer> id = font.getfamilyoffont(<integer> fam)
```

10.3.14.9 [set|get]interaction

The engine can in one of four modes:

VALUE	mode	MEANING
0	batch	omits all stops and omits terminal output
1	nonstop	omits all stops
2	scroll	omits error stops
3	errorstop	stops at every opportunity to interact

The mode can be queried and set with:

```
<integer> i = tex.getinteraction()  
tex.setinteraction(<integer> i)
```

10.3.14.10 runtoks and quittoks

Because of the fact that T_EX is in a complex dance of expanding, dealing with fonts, typesetting paragraphs, messing around with boxes, building pages, and so on, you cannot easily run a nested T_EX run (read nested main loop). However, there is an option to force a local run with `runtoks`. The content of the given token list register gets expanded locally after which we return to where we triggered this expansion, at the LUA end. Instead a function can get passed that does some work. You have to make sure that at the end T_EX is in a sane state and this is not always trivial. A more complex mechanism would complicate T_EX itself (and probably also harm performance) so this simple local expansion loop has to do.

```
tex.runtoks(<token register>)  
tex.runtoks(<lua function>)  
tex.runtoks(<macro name>)  
tex.runtoks(<register name>)
```

When the `\tracingnesting` parameter is set to a value larger than 2 some information is reported about the state of the local loop. The return value indicates an error:

VALUE	meaning
0	no error
1	bad register number
2	unknown macro or register name
3	macro is unsuitable for runtoks (has arguments)

This function has two optional arguments in case a token register is passed:



```
tex.runtoks(<token register>,force,grouped)
```

Inside for instance an `\edef` the `runtoks` function behaves (at least tries to) like it were an `\the`. This prevents unwanted side effects: normally in such a definition tokens remain tokens and (for instance) characters don't become nodes. With the second argument you can force the local main loop, no matter what. The third argument adds a level of grouping.

You can quit the local loop with `\endlocalcontrol` or from the LUA end with `tex.quittoks`. In that case you end one level up! Of course in the end that can mean that you arrive at the main level in which case an extra end will trigger a redundancy warning (not an abort!).

10.3.14.11 **forcehmode**

An example of a (possible error triggering) complication is that \TeX expects to be in some state, say horizontal mode, and you have to make sure it is when you start feeding back something from LUA into \TeX . Normally a user will not run into issues but when you start writing tokens or nodes or have a nested run there can be situations that you need to run `forcehmode`. There is no recipe for this and intercepting possible cases would weaken $\text{LUA}\TeX$'s flexibility.

10.3.14.12 **hashtokens**

```
for i,v in pairs (tex.hashtokens()) do ... end
```

Returns a list of names. This can be useful for debugging, but note that this also reports control sequences that may be unreachable at this moment due to local redefinitions: it is strictly a dump of the hash table. You can use `token.create` to inspect properties, for instance when the command key in a created table equals 123, you have the `cmdname` value `undefined_cs`.

10.3.14.13 **definefont**

```
tex.definefont(<string> csname, <number> fontid)
tex.definefont(<boolean> global, <string> csname, <number> fontid)
```

Associates `csname` with the internal font number `fontid`. The definition is global if (and only if) `global` is specified and true (the setting of `globaldefs` is not taken into account).

10.3.15 **Functions for dealing with primitives**

10.3.15.1 **enableprimitives**

```
tex.enableprimitives(<string> prefix, <table> primitive names)
```

This function accepts a prefix string and an array of primitive names. For each combination of 'prefix' and 'name', the `tex.enableprimitives` first verifies that 'name' is an actual primitive (it must be returned by one of the `tex.extraprimitives` calls explained below, or part of $\text{T}\text{E}\text{X}82$, or `\directlua`). If it is not, `tex.enableprimitives` does nothing and skips to the next pair.



But if it is, then it will construct a `csname` variable by concatenating the ‘prefix’ and ‘name’, unless the ‘prefix’ is already the actual prefix of ‘name’. In the latter case, it will discard the ‘prefix’, and just use ‘name’.

Then it will check for the existence of the constructed `csname`. If the `csname` is currently undefined (note: that is not the same as `\relax`), it will globally define the `csname` to have the meaning: run code belonging to the primitive ‘name’. If for some reason the `csname` is already defined, it does nothing and tries the next pair.

An example:

```
tex.enableprimitives('LuaTeX', {'formatname'})
```

will define `\LuaTeXformatname` with the same intrinsic meaning as the documented primitive `\formatname`, provided that the control sequences `\LuaTeXformatname` is currently undefined.

When `LUATEX` is run with `--ini` only the `TEX82` primitives and `\directlua` are available, so no extra primitives **at all**.

If you want to have all the new functionality available using their default names, as it is now, you will have to add

```
\ifx\directlua\undefined \else
  \directlua {tex.enableprimitives('',tex.extraprimitives ())}
\fi
```

near the beginning of your format generation file. Or you can choose different prefixes for different subsets, as you see fit.

Calling some form of `tex.enableprimitives` is highly important though, because if you do not, you will end up with a `TEX82`-lookalike that can run `LUA` code but not do much else. The defined `csnames` are (of course) saved in the format and will be available at runtime.

10.3.15.2 extraprimitives

```
<table> t = tex.extraprimitives(<string> s, ...)
```

This function returns a list of the primitives that originate from the engine(s) given by the requested string value(s). The possible values and their (current) return values are given in the following table. In addition the somewhat special primitives ‘\ ’, ‘\/' and ‘-’ are defined.

NAME	VALUES
tex	vskip vsize unhcopy output unskip unvbox boxmaxdepth muskipdef string toks- def floatingpenalty righthyphenmin voffset escapechar topmark splitfirst- mark vsplit everydisplay badness xleaders textfont showlists language math- choice topskip abovedisplayshortskip underline tracinglostchars pagefill- stretch unvcopy splitbotmark finalhyphendemerits atopwithdelims pretoler- ance dp setlanguage ht mathchardef nulldelimiterspace or wd pagegoal advance chardef catcode mathchar scriptscriptfont mathcode leftskip pageshrink pagefilstretch delcode fontname brokenpenalty lastkern belowdisplayshort- skip tolerance mathopen exhyphenpenalty maxdepth futurelet abovewithdelims



hangindent lastskip delimiterfactor everyjob xspaceskip globaldefs every-
 tab scriptfont delimiter afterassignment firstmark lineskiplimit lineskip
 def fam day iffalse textstyle end mag box belowdisplayskip ifx let errmessage
 exhyphenchar hss expandafter hfilneg the displaywidth pagedeepth looseness
 leaders vss ifhmode botmark displaystyle accent ifmmode meaning abovedis-
 playskip medmuskip emergencystretch rightskip mathclose hangafter hoff-
 set aftergroup cleaders romannumeral hbadness mathbin showboxbreadth ifv-
 mode jobname vbadness patterns nonstopmode errhelp predisplayskip end-
 linechar mathinner lastbox showboxdepth postdisplayskip mathrel ordlim-
 its holdinginserts radical mathord pagetotal everycr adjdemerits halign de-
 faultskewchar errorcontextlines splitmaxdepth ifcase errorstopmode noin-
 dent tracingmacros moveright predisplaysize tracingrestores message ifh-
 box deadcycles interlinepenalty mathpunct lccode displayindent nonscript
 everyhbox global penalty tracingcommands everymath nolimits noalign in-
 putlineno parskip indent dimendef widowpenalty ifvbox above spaceskip mid-
 dle displaylimits pausing everyvbox iftrue moveleft mathop endcsname di-
 men ifcat clubpenalty splittopskip doublehyphendemerits ifdim limits ifeof
 ignorespaces insert delimitershortfall ifodd insertpenalties tracingpages
 vadjust tracingonline count char begingroup tracingparagraphs hyphenation
 hfuzz legno hyphenchar hyphenpenalty vcenter hfil thickmuskip maxdeadcy-
 cles skewchar hbox overfullrule else hsize atop thinmuskip spacefactor input
 hrule left eqno parfillskip font valign kern relax prevdepth shipout batch-
 mode right setbox baselineskip mskip endgroup uchyph binoppenalty endinput
 omit overwithdelims newlinechar vfilneg time skip span prevgraf over vbox
 tracingstats year defaultthyphenchar nullfont muskip toks multiply tracin-
 goutput parindent displaywidowpenalty unhbox lefthyphenmin vtop mathaccent
 showthe vfuzz overline unkern crcr cr fi par showbox show vfill vfil raise
 lower dump skipdef countdef scrollmode unpenalty outer long mkern script-
 style scriptscriptstyle copy openin closein hskip hfill discretionary xdef
 gdef edef uccode sfcode uppercase lowercase orelse ifvoid ifinner ifnum if
 number lastpenalty pagefillllstretch pagestretch vrule parshape read noex-
 pand mark csname fontdimen divide scriptspace mathsurround shownodedetails
 relpenalty outputpenalty month linepenalty everypar tabskip

core

directlua

etex

unless botmarks currentifttype pagediscards mutoglu displaywidowpenalties
 fontcharic fontchardp fontcharht fontcharwd widowpenalties tracingifs if-
 fontchar splitdiscards eTeXVersion protected topmarks showgroups glueexpr
 splitfirstmarks predisplayskip direction gluetomu everyeof eTeXversion scant-
 oken clubpenalties savingvdiscards splitbotmarks showtokens tracingassigns
 dimexpr parshapedimen readline eTeXminorversion tracingscantokens trac-
 ingnesting glueshrinkorder ifdefined lastlinefit gluestretchorder tracing-
 groups currentifbranch firstmarks lastnodetype marks currentgrouplevel in-
 terlinepenalties interactionmode muexpr unexpanded savinghyphcodes ifcsname
 parshapeindent numexpr showifs parshapelength eTeXrevision gluestretch
 glueshrink currentiflevel currentgroupstype detokenize



luatex Umathcloseopspacing boxxoffset Umathordpunctspacing Udelimiterunder bound-
 ary boxorientation mathsurroundmode Uskewedwithdelims Umathopenpunctspac-
 ing boxyoffset iffrozen luabytecodecall mathsurroundskip endllocalcontrol
 scantexttokens Umathbinclasespacing toksapp rightghost Umathlimitbelow-
 bgap Umathopeninnerspacing textdirection tokspre Umathnolimitsubfactor
 Uoverdelimiter Umathpunctpunctspacing Umathclosepunctspacing ifprotected
 mathdisplayskipmode mathrulesfam Umathrelordspacing Umathsupbottommin ad-
 justspacingstep Umathlimitbelowkern internalcodesmode Umathstackdenom-
 down localrightbox Umathfractionrule Umathcharfam Umathcloseinnerspac-
 ing Umathopenrelspacing Uextentensible Umathsupsubbottommax leftmarginkern
 Umathcloserelspacing linedirection ifincsname Umathcharnum boxtotal Umath-
 innerordspacing ifnumval beginlocalcontrol luabytecode formatname boxdi-
 rection boxxmove unletprotected boxymove Umathsubtopmax mathscriptsmode
 Umathsubsupshiftdown Umathopbinspacing Umathordbinspacing Umathrelopspac-
 ing Umathopenbinspacing Umathoverdelimiterbgap localleftbox alignmark Uun-
 derdelimiter hyphenationmin Umathclosebinspacing Umathcodenum luafunction
 compoundhyphenmode Umathpunctopenspacing luacopyinputnodes Umathconnec-
 toroverlapmin letprotected crampedscriptscriptstyle csstring Umathradic-
 aldegreeafter luatexversion Umathfractionnumup rightmarginkern Umath-
 opclosespacing mathrulesmode explicithyphenpenalty Umathordclosespacing
 Umathoverdelimitervgap etokspre expanded Udelcode wordboundary immediate-
 assigned shapemode attribute preexhyphenchar Umathsubshiftdown matheqnogap-
 step Umathpunctrelspacing Umathradicaldegreeraise fixupboxesmode mathno-
 limitsmode adjustspacing Umathsupshiftdrop Umathcharslot Umathcloseclos-
 espacing luatexrevision insertht localinterlinepenalty aftergrouped ex-
 plicitdiscretionary Umathchar ifcstok Udelimiterover Ustack Umathcode math-
 delimitersmode Udelcodenum gtoksapp boxattr Umathaxis Umathfractionnumv-
 gap gtokspre mathflattenmode Umathskewedfractionhgap Umathrelclosespacing
 Umathpunctbinspacing Ustopdisplaymath quitvmode adjustspacingstretch
 crampedscriptstyle letcharcode hyphenationbounds crampedtextstyle Umath-
 binrelspacing Umathopordspacing attributedef mathdirection Umathor-
 dordspacing ifusercmd Umathskewedfractionvgap Umathopenordspacing math-
 italicsmode outputbox Umathcloseordspacing Umathnolimitsupfactor Ustop-
 math aligntab prehyphenchar luafunctioncall Umathpunctopspacing breakaf-
 terdirmode Umathsubsupvgap luaescapestring adjustspacingshrink beginc-
 sname Umathradicalrule postexhyphenchar Umathradicaldegreebefore Umath-
 binbinspacing hjcode noboundary Umathbinopspacing xtoksapp Ustartdisplay-
 math savecatcodetable matholdmode mathscriptboxmode Uroot Unosuperscript
 Umathoperatorsizetokspre Uradical mathstyle Umathopopenspacing Umathor-
 dopenspacing automatichyphenpenalty Umathinnerrelspacing clearmarks Umath-
 overbarvgap fontid Umathopenopenspacing immediateassignment Umathover-
 barrule setfontid crampeddisplaystyle ifabsdim Umathlimitabovebgap Umath-
 charclass Umathstackvgap Umathinneropspacing Umathrelbinspacing Umath-
 closeopenspacing ifcondition futureexpandis frozen initcatcodetable nokerns
 ifabsnum letfrozen hpack luaodef tracingfonts nospaces Umathrelopenspacing
 ifchkdim Umathlimitabovekern glyphdatafield Udelimiter iftok nohrule math-



rulethicknessmode localbrokenpenalty lastnodesubtype Umathfractiondel-
size exceptionpenalty automaticdiscretionary gleaders Umathinnerbinspacing
ifchknum noligs ignorepars hyphenpenalty Ustartmath futureexpand ifdim-
val ifcmpdim automatichyphenmode prebinoppenalty Usubscript Umathcharnumdef
rcode tpack unletfrozen ifcmpnum mathpenaltiesmode mathscriptcharmode
vpack Umathaccent firstvalidlanguage catcodetable protrusionboundary Umath-
spaceafterscript predisplayspacefactor Umathinneropenspacing Uskewed pxd-
imen Umathordopspacing Umathopenopspacing Umathunderdelimitervgap Umath-
underdelimiterbgap Umathunderbarvgap Umathunderbarrule Umathunderbark-
ern Umathsupshiftup Umathsubshiftdrop Umathstacknumup Umathrelrelspacing
Umathrelpunctspacing Umathrelinnerspacing Umathradicalvgap Umathradicalk-
ern Umathquad Umathpunctordspacing Umathpunctinnerspacing Umathpunctclos-
espacing Umathoverbarkern Umathordrelspacing Umathordinnerspacing Umath-
oprelspacing Umathoppunctspacing Umathopopspacing Umathopinnerspacing
Umathopenclosespacing Umathlimitbelowvgap Umathlimitabovevgap Umathinner-
punctspacing Umathinnerinnerspacing Umathinnerclosespacing Umathfraction-
denomvgap Umathfractiondenomdown Umathbinpunctspacing Umathbinordspacing
Umathbinopenspacing Umathbininnerspacing Uunosubscript Usuperscript pardi-
rection Umathchardef glet Uvextensible Uright Umiddle Uleft posthyphenchar
lrcode efcodes Uchar luatexbanner novrule futureexpandisap lastnamedcs etok-
sapp leftghost protrudechars prerelpenalty

Note that `lua` does not contain `directlua`, as that is considered to be a core primitive, along with all the \TeX 82 primitives, so it is part of the list that is returned from 'core'.

Running `tex.extraprimitives` will give you the complete list of primitives -ini startup. It is exactly equivalent to `tex.extraprimitives("etex", "luatex")`.

10.3.15.3 primitives

```
<table> t = tex.primitives()
```

This function returns a list of all primitives that \LaTeX knows about.

10.3.16 Core functionality interfaces

10.3.16.1 badness

```
<number> b = tex.badness(<number> t, <number> s)
```

This helper function is useful during linebreak calculations. `t` and `s` are scaled values; the function returns the badness for when total `t` is supposed to be made from amounts that sum to `s`. The returned number is a reasonable approximation of $100(t/s)^3$;

10.3.16.2 tex.resetparagraph

This function resets the parameters that \TeX normally resets when a new paragraph is seen.



10.3.16.3 linebreak

```
local <node> nodelist, <table> info =  
    tex.linebreak(<node> listhead, <table> parameters)
```

The understood parameters are as follows:

NAME	TYPE	EXPLANATION
pardir	string	
pretolerance	number	
tracingparagraphs	number	
tolerance	number	
looseness	number	
hyphenpenalty	number	
exhyphenpenalty	number	
pdfadjustspacing	number	
adjdemerits	number	
pdfprotrudechars	number	
linepenalty	number	
lastlinefit	number	
doublehyphendemerits	number	
finalhyphendemerits	number	
hangafter	number	
interlinepenalty	number or table	if a table, then it is an array like <code>\interlinepenalties</code>
clubpenalty	number or table	if a table, then it is an array like <code>\clubpenalties</code>
widowpenalty	number or table	if a table, then it is an array like <code>\widowpenalties</code>
brokenpenalty	number	
emergencystretch	number	in scaled points
hangindent	number	in scaled points
hsize	number	in scaled points
leftskip	glue_spec node	
rightskip	glue_spec node	
parshape	table	

Note that there is no interface for `\displaywidowpenalties`, you have to pass the right choice for `widowpenalties` yourself.

It is your own job to make sure that `listhead` is a proper paragraph list: this function does not add any nodes to it. To be exact, if you want to replace the core line breaking, you may have to do the following (when you are not actually working in the `pre_linebreak_filter` or `linebreak_filter` callbacks, or when the original list starting at `listhead` was generated in horizontal mode):

- ▶ add an ‘indent box’ and perhaps a `local_par` node at the start (only if you need them)
- ▶ replace any found final glue by an infinite penalty (or add such a penalty, if the last node is not a glue)
- ▶ add a glue node for the `\parfillskip` after that penalty node
- ▶ make sure all the prev pointers are OK



The result is a node list, it still needs to be vpacked if you want to assign it to a `\vbox`. The returned info table contains four values that are all numbers:

NAME	EXPLANATION
<code>prevdepth</code>	depth of the last line in the broken paragraph
<code>prevgraf</code>	number of lines in the broken paragraph
<code>looseness</code>	the actual looseness value in the broken paragraph
<code>demerits</code>	the total demerits of the chosen solution

Note there are a few things you cannot interface using this function: You cannot influence font expansion other than via `pdfadjustspacing`, because the settings for that take place elsewhere. The same is true for `hbadness` and `hfuzz` etc. All these are in the `hpack` routine, and that fetches its own variables via globals.

10.3.16.4 `shipout`

```
tex.shipout(<number> n)
```

Ships out box number `n` to the output file, and clears the box register.

10.3.16.5 `getpagestate`

This helper reports the current page state: `empty`, `box_there` or `inserts_only` as integer value.

10.3.16.6 `getlocallevel`

This integer reports the current level of the local loop. It's only useful for debugging and the (relative state) numbers can change with the implementation.

10.3.17 Randomizers

For practical reasons L^AT_EX has its own random number generator. The original LUA random function is available as `tex.lua_math_random`. You can initialize with a new seed with `init_rand` (`lua_math_randomseed` is equivalent to this one).

There are three generators: `normal_rand` (no argument is used), `uniform_rand` (takes a number that will get rounded before being used) and `uniformdeviate` which behaves like the primitive and expects a scaled integer, so

```
tex.print(tex.uniformdeviate(65536)/65536)
```

will give a random number between zero and one.

10.3.18 Functions related to `synctex`

The next helpers only make sense when you implement your own `synctex` logic. Keep in mind that the library used in editors assumes a certain logic and is geared for plain and L^AT_EX, so after a decade users expect a certain behaviour.



NAME	EXPLANATION
set_synctex_mode	0 is the default and used normal synctex logic, 1 uses the values set by the next helpers while 2 also sets these for glyph nodes; 3 sets glyphs and glue and 4 sets only glyphs
set_synctex_tag	set the current tag (file) value (obeys save stack)
set_synctex_line	set the current line value (obeys save stack)
set_synctex_no_files	disable synctex file logging
get_synctex_mode	returns the current mode (for values see above)
get_synctex_tag	get the currently set value of tag (file)
get_synctex_line	get the currently set value of line
force_synctex_tag	overload the tag (file) value (0 resets)
force_synctex_line	overload the line value (0 resets)

The last one is somewhat special. Due to the way files are registered in SYNCT_EX we need to explicitly disable that feature if we provide our own alternative if we want to avoid that overhead. Passing a value of 1 disables registering.

10.4 The texconfig table

This is a table that is created empty. A startup LUA script could fill this table with a number of settings that are read out by the executable after loading and executing the startup file.

KEY	TYPE	DEFAULT
max_strings	number	100000
strings_free	number	100
nest_size	number	50
max_in_open	number	100
param_size	number	60
save_size	number	5000
stack_size	number	500
expand_depth	number	1000
function_size	number	0
error_line	number	79
half_error_line	number	50
hash_extra	number	0
formatname	string	
jobname	string	

If no format name or jobname is given on the command line, the related keys will be tested first instead of simply quitting.



10.5 The texio library

This library takes care of the low-level I/O interface: writing to the log file and/or console.

10.5.1 write

```
texio.write(<string> target, <string> s, ...)
texio.write(<string> s, ...)
```

Without the target argument, writes all given strings to the same location(s) \TeX writes messages to at this moment. If `\batchmode` is in effect, it writes only to the log, otherwise it writes to the log and the terminal. The optional target can be one of three possibilities: `term`, `log` or `term and log`.

Note: If several strings are given, and if the first of these strings is or might be one of the targets above, the target must be specified explicitly to prevent LUA from interpreting the first string as the target.

10.5.2 write_nl

```
texio.write_nl(<string> target, <string> s, ...)
texio.write_nl(<string> s, ...)
```

This function behaves like `texio.write`, but make sure that the given strings will appear at the beginning of a new line. You can pass a single empty string if you only want to move to the next line.

10.5.3 setescape

You can disable ^^ escaping of control characters by passing a value of zero.

10.5.4 closeinput

This function that should be used with care. It acts as `\endinput` but at the LUA end. You can use it to (sort of) force a jump back to \TeX . Normally a LUA will just collect prints and at the end bump an input level and flush these prints. This function can help you stay at the current level but you need to know what you're doing (or more precise: what \TeX is doing with input).

10.6 The token library

10.6.1 The scanner

The token library provides means to intercept the input and deal with it at the LUA level. The library provides a basic scanner infrastructure that can be used to write macros that accept a wide range of arguments. This interface is on purpose kept general and as performance is



quite ok. One can build additional parsers without too much overhead. It's up to macro package writers to see how they can benefit from this as the main principle behind L^AT_EX is to provide a minimal set of tools and no solutions. The scanner functions are probably the most intriguing.

FUNCTION	ARGUMENT	RESULT
<code>scan_keyword</code>	string	returns true if the given keyword is gobbled; as with the regular T _E X keyword scanner this is case insensitive (and ASCII based)
<code>scan_keywordcs</code>	string	returns true if the given keyword is gobbled; this variant is case sensitive and also suitable for UTF8
<code>scan_int</code>		returns an integer
<code>scan_real</code>		returns a number from e.g. 1, 1.1, .1 with optional collapsed signs
<code>scan_float</code>		returns a number from e.g. 1, 1.1, .1, 1.1E10, , .1e-10 with optional collapsed signs
<code>scan_dimen</code>	infinity, mu-units	returns a number representing a dimension and or two numbers being the filler and order
<code>scan_glue</code>	mu-units	returns a glue spec node
<code>scan_toks</code>	definer, expand	returns a table of tokens tokens
<code>scan_code</code>	bitset	returns a character if its category is in the given bitset (representing catcodes)
<code>scan_string</code>		returns a string given between {}, as \macro or as sequence of characters with catcode 11 or 12
<code>scan_argument</code>		this one is similar to <code>scanstring</code> but also accepts a \cs (which then get expanded)
<code>scan_word</code>		returns a sequence of characters with catcode 11 or 12 as string
<code>scan_csname</code>		returns foo after scanning \foo
<code>scan_list</code>		picks up a box specification and returns a [h v]list node

The scanners can be considered stable apart from the one scanning for a token. The `scan_code` function takes an optional number, the keyword function a normal LUA string. The `infinity` boolean signals that we also permit fill as dimension and the `mu-units` flags the scanner that we expect math units. When scanning tokens we can indicate that we are defining a macro, in which case the result will also provide information about what arguments are expected and in the result this is separated from the meaning by a separator token. The `expand` flag determines if the list will be expanded.

The string scanner scans for something between curly braces and expands on the way, or when it sees a control sequence it will return its meaning. Otherwise it will scan characters with catcode letter or other. So, given the following definition:

```
\def\bar{bar}
\def\foo{foo-\bar}
```

we get:

NAME	RESULT
<code>\directlua{token.scan_string()}{foo}</code>	full expansion



<code>\directlua{token.scan_string()}foo</code>	letters and others
<code>\directlua{token.scan_string()}\foo</code>	meaning

The `\foo` case only gives the meaning, but one can pass an already expanded definition (`\edef'd`). In the case of the braced variant one can of course use the `\detokenize` and `\unexpanded` primitives since there we do expand.

The `scan_word` scanner can be used to implement for instance a number scanner:

```
function token.scan_number(base)
    return tonumber(token.scan_word(),base)
end
```

This scanner accepts any valid LUA number so it is a way to pick up floats in the input.

You can use the LUA interface as follows:

```
\directlua {
    function mymacro(n)
        ...
    end
}

\def\mymacro#1{%
    \directlua {
        mymacro(\number\dimexpr#1)
    }%
}

\mymacro{12pt}
\mymacro{\dimen0}
```

You can also do this:

```
\directlua {
    function mymacro()
        local d = token.scan_dimen()
        ...
    end
}

\def\mymacro{%
    \directlua {
        mymacro()
    }%
}

\mymacro 12pt
\mymacro \dimen0
```



It is quite clear from looking at the code what the first method needs as argument(s). For the second method you need to look at the LUA code to see what gets picked up. Instead of passing from T_EX to LUA we let LUA fetch from the input stream.

In the first case the input is tokenized and then turned into a string, then it is passed to LUA where it gets interpreted. In the second case only a function call gets interpreted but then the input is picked up by explicitly calling the scanner functions. These return proper LUA variables so no further conversion has to be done. This is more efficient but in practice (given what T_EX has to do) this effect should not be overestimated. For numbers and dimensions it saves a bit but for passing strings conversion to and from tokens has to be done anyway (although we can probably speed up the process in later versions if needed).

10.6.2 Picking up one token

The scanners look for a sequence. When you want to pick up one token from the input you use `get_next`. This creates a token with the (low level) properties as discussed next. This token is just the next one. If you want to enforce expansion first you can use `scan_token`. Internally tokens are characterized by a number that packs a lot of information. In order to access the bits of information a token is wrapped in a userdata object.

The `expand` function will trigger expansion of the next token in the input. This can be quite unpredictable but when you call it you probably know enough about T_EX not to be too worried about that. It basically is a call to the internal `expand` related function.

10.6.3 Creating tokens

The creator function can be used as follows:

```
local t = token.create("relax")
```

This gives back a token object that has the properties of the `\relax` primitive. The possible properties of tokens are:

NAME	EXPLANATION
<code>command</code>	a number representing the internal command number
<code>cmdname</code>	the type of the command (for instance the catcode in case of a character or the classifier that determines the internal treatment)
<code>csname</code>	the associated control sequence (if applicable)
<code>id</code>	the unique id of the token
<code>tok</code>	the full token number as stored in T _E X
<code>active</code>	a boolean indicating the active state of the token
<code>expandable</code>	a boolean indicating if the token (macro) is expandable
<code>protected</code>	a boolean indicating if the token (macro) is protected
<code>mode</code>	a number either representing a character or another entity
<code>index</code>	a number running from 0x0000 upto 0xFFFF indicating a T _E X register index

Alternatively you can use a getter `get_<fieldname>` to access a property of a token.



The numbers that represent a catcode are the same as in T_EX itself, so using this information assumes that you know a bit about T_EX's internals. The other numbers and names are used consistently but are not frozen. So, when you use them for comparing you can best query a known primitive or character first to see the values.

You can ask for a list of commands:

```
local t = token.commands()
```

The id of a token class can be queried as follows:

```
local id = token.command_id("math_shift")
```

If you really know what you're doing you can create character tokens by not passing a string but a number:

```
local letter_x = token.create(string.byte("x"))
local other_x = token.create(string.byte("x"), 12)
```

Passing weird numbers can give side effects so don't expect too much help with that. As said, you need to know what you're doing. The best way to explore the way these internals work is to just look at how primitives or macros or `\chardef`'d commands are tokenized. Just create a known one and inspect its fields. A variant that ignores the current catcode table is:

```
local whatever = token.new(123, 12)
```

You can test if a control sequence is defined with `is_defined`, which accepts a string and returns a boolean:

```
local okay = token.is_defined("foo")
```

The largest character possible is returned by `biggest_char`, just in case you need to know that boundary condition.

10.6.4 Macros

The `set_macro` function can get upto 4 arguments:

```
set_macro("csname", "content")
set_macro("csname", "content", "global")
set_macro("csname")
```

You can pass a catcodetable identifier as first argument:

```
set_macro(catcodetable, "csname", "content")
set_macro(catcodetable, "csname", "content", "global")
set_macro(catcodetable, "csname")
```

The results are like:

```
\def\csname{content}
```



```
\gdef\csname{content}  
\def\csname{}
```

The `get_macro` function can be used to get the content of a macro while the `get_meaning` function gives the meaning including the argument specification (as usual in \TeX separated by `->`).

The `set_char` function can be used to do a `\chardef` at the Lua end, where invalid assignments are silently ignored:

```
set_char("csname",number)  
set_char("csname",number,"global")
```

A special one is the following:

```
set_lua("mycode",id)  
set_lua("mycode",id,"global","protected")
```

This creates a token that refers to a Lua function with an entry in the table that you can access with `lua.get_functions_table`. It is the companion to `\luadef`. When the first (and only) argument is true the size will preset to the value of `texconfig.function_size`.

10.6.5 Pushing back

There is a (for now) experimental putter:

```
local t1 = token.get_next()  
local t2 = token.get_next()  
local t3 = token.get_next()  
local t4 = token.get_next()  
-- watch out, we flush in sequence  
token.put_next { t1, t2 }  
-- but this one gets pushed in front  
token.put_next ( t3, t4 )
```

When we scan `wxyz!` we get `yzwx!` back. The argument is either a table with tokens or a list of tokens. The `token.expand` function will trigger expansion but what happens really depends on what you're doing where.

10.6.6 Nota bene

When scanning for the next token you need to keep in mind that we're not scanning like \TeX does: expanding, changing modes and doing things as it goes. When we scan with Lua we just pick up tokens. Say that we have:

```
\bar
```

but `\bar` is undefined. Normally \TeX will then issue an error message. However, when we have:

```
\def\foo{\bar}
```



We get no error, unless we expand `\foo` while `\bar` is still undefined. What happens is that as soon as \TeX sees an undefined macro it will create a hash entry and when later it gets defined that entry will be reused. So, `\bar` really exists but can be in an undefined state.

```
bar : bar
foo : foo
myfirstbar :
```

This was entered as:

```
bar      : \directlua{tex.print(token.scan_csname())}\bar
foo      : \directlua{tex.print(token.scan_csname())}\foo
myfirstbar : \directlua{tex.print(token.scan_csname())}\myfirstbar
```

The reason that you see `bar` reported and not `myfirstbar` is that `\bar` was already used in a previous paragraph.

If we now say:

```
\def\foo{}
```

we get:

```
bar : bar
foo : foo
myfirstbar :
```

And if we say

```
\def\foo{\bar}
```

we get:

```
bar : bar
foo : foo
myfirstbar :
```

When scanning from `LUA` we are not in a mode that defines (undefined) macros at all. There we just get the real primitive undefined macro token.

```
760729 536941995
757810 536941995
757625 536941995
```

This was generated with:

```
\directlua{local t = token.get_next() tex.print(t.id.." "..t.tok)}\myfirstbar
\directlua{local t = token.get_next() tex.print(t.id.." "..t.tok)}\mysecondbar
\directlua{local t = token.get_next() tex.print(t.id.." "..t.tok)}\mythirdbar
```

So, we do get a unique token because after all we need some kind of `LUA` object that can be used and garbage collected, but it is basically the same one, representing an undefined control sequence.



11 The METAPOST library `mplib`

11.1 Process management

The METAPOST library interface registers itself in the table `mplib`. It is based on MPLIB version 3.01.

11.1.1 new

To create a new METAPOST instance, call

```
<mpinstance> mp = mplib.new({...})
```

This creates the `mp` instance object. The argument hash can have a number of different fields, as follows:

NAME	TYPE	DESCRIPTION	DEFAULT
<code>error_line</code>	number	error line width	79
<code>print_line</code>	number	line length in ps output	100
<code>random_seed</code>	number	the initial random seed	variable
<code>math_mode</code>	string	the number system to use: scaled, double or decimal	scaled
<code>interaction</code>	string	the interaction mode: batch, nonstop, scroll or errorstop	errorstop
<code>job_name</code>	string	a compatibility value	mpout
<code>find_file</code>	function	a function to find files	only local files
<code>utf8</code>	boolean	permit characters in the range 128 upto 255 to be part of names	false

The binary mode is no longer available in the L^AT_EX version of MPLIB. It offers no real advantage and brings a ton of extra libraries with platform specific properties that we can now avoid. We might introduce a high resolution scaled variant at some point but only when it pays of performance wise.

The `find_file` function should be of this form:

```
<string> found = finder (<string> name, <string> mode, <string> type)
```

with:

NAME	THE REQUESTED FILE
<code>mode</code>	the file mode: r or w
<code>type</code>	the kind of file, one of: mp, tfm, map, pfb, enc

Return either the full path name of the found file, or `nil` if the file cannot be found.

Note that the new version of MPLIB no longer uses binary mem files, so the way to preload a set of macros is simply to start off with an input command in the first execute call.



When you are processing a snippet of text starting with `btex` and ending with either `etex` or `verbatimtex`, the METAPOST `texscriptmode` parameter controls how spaces and newlines get honoured. The default value is 1. Possible values are:

NAME	MEANING
0	no newlines
1	newlines in <code>verbatimtex</code>
2	newlines in <code>verbatimtex</code> and <code>etex</code>
3	no leading and trailing strip in <code>verbatimtex</code>
4	no leading and trailing strip in <code>verbatimtex</code> and <code>btex</code>

That way the LUA handler (assigned to `make_text`) can do what it likes. An `etex` has to be followed by a space or `;` or be at the end of a line and preceded by a space or at the beginning of a line.

11.1.2 statistics

You can request statistics with:

```
<table> stats = mp:statistics()
```

This function returns the vital statistics for an MPLIB instance. There are four fields, giving the maximum number of used items in each of four allocated object classes:

FIELD	TYPE	EXPLANATION
<code>main_memory</code>	number	memory size
<code>hash_size</code>	number	hash size
<code>param_size</code>	number	simultaneous macro parameters
<code>max_in_open</code>	number	input file nesting levels

Note that in the new version of MPLIB, this is informational only. The objects are all allocated dynamically, so there is no chance of running out of space unless the available system memory is exhausted.

11.1.3 execute

You can ask the METAPOST interpreter to run a chunk of code by calling

```
<table> rettable = execute(mp,"metapost code")
```

for various bits of METAPOST language input. Be sure to check the `rettable.status` (see below) because when a fatal METAPOST error occurs the MPLIB instance will become unusable thereafter.

Generally speaking, it is best to keep your chunks small, but beware that all chunks have to obey proper syntax, like each of them is a small file. For instance, you cannot split a single statement over multiple chunks.



In contrast with the normal stand alone mpost command, there is *no* implied ‘input’ at the start of the first chunk.

11.1.4 finish

```
<table> rettable = finish(mp)
```

If for some reason you want to stop using an MPLIB instance while processing is not yet actually done, you can call `finish`. Eventually, used memory will be freed and open files will be closed by the LUA garbage collector, but an explicit `finish` is the only way to capture the final part of the output streams.

11.2 The end result

The return value of `execute` and `finish` is a table with a few possible keys (only `status` is always guaranteed to be present).

FIELD	TYPE	EXPLANATION
log	string	output to the ‘log’ stream
term	string	output to the ‘term’ stream
error	string	output to the ‘error’ stream (only used for ‘out of memory’)
status	number	the return value: 0 = good, 1 = warning, 2 = errors, 3 = fatal error
fig	table	an array of generated figures (if any)

When `status` equals 3, you should stop using this MPLIB instance immediately, it is no longer capable of processing input.

If it is present, each of the entries in the `fig` array is a userdata representing a figure object, and each of those has a number of object methods you can call:

FIELD	TYPE	EXPLANATION
boundingbox	function	returns the bounding box, as an array of 4 values
postscript	function	returns a string that is the ps output of the fig. this function accepts two optional integer arguments for specifying the values of prologues (first argument) and procset (second argument)
svg	function	returns a string that is the svg output of the fig. This function accepts an optional integer argument for specifying the value of prologues
objects	function	returns the actual array of graphic objects in this fig
copy_objects	function	returns a deep copy of the array of graphic objects in this fig
filename	function	the filename this fig’s POSTSCRIPT output would have written to in stand alone mode
width	function	the fontcharwd value
height	function	the fontcharht value
depth	function	the fontchardp value
italcorr	function	the fontcharit value
charcode	function	the (rounded) charcode value



Note: you can call `fig:objects()` only once for any one `fig` object!

When the boundingbox represents a ‘negated rectangle’, i.e. when the first set of coordinates is larger than the second set, the picture is empty.

Graphical objects come in various types that each has a different list of accessible values. The types are: `fill`, `outline`, `text`, `start_clip`, `stop_clip`, `start_bounds`, `stop_bounds`, `special`.

There is a helper function (`mplib.fields(obj)`) to get the list of accessible values for a particular object, but you can just as easily use the tables given below.

All graphical objects have a field type that gives the object type as a string value; it is not explicit mentioned in the following tables. In the following, numbers are POSTSCRIPT points represented as a floating point number, unless stated otherwise. Field values that are of type `table` are explained in the next section.

11.2.1 fill

FIELD	TYPE	EXPLANATION
<code>path</code>	<code>table</code>	the list of knots
<code>htap</code>	<code>table</code>	the list of knots for the reversed trajectory
<code>pen</code>	<code>table</code>	knots of the pen
<code>color</code>	<code>table</code>	the object’s color
<code>linejoin</code>	<code>number</code>	line join style (bare number)
<code>miterlimit</code>	<code>number</code>	miterlimit
<code>prescript</code>	<code>string</code>	the prescript text
<code>postscript</code>	<code>string</code>	the postscript text

The entries `htap` and `pen` are optional.

11.2.2 outline

FIELD	TYPE	EXPLANATION
<code>path</code>	<code>table</code>	the list of knots
<code>pen</code>	<code>table</code>	knots of the pen
<code>color</code>	<code>table</code>	the object’s color
<code>linejoin</code>	<code>number</code>	line join style (bare number)
<code>miterlimit</code>	<code>number</code>	miterlimit
<code>linecap</code>	<code>number</code>	line cap style (bare number)
<code>dash</code>	<code>table</code>	representation of a dash list
<code>prescript</code>	<code>string</code>	the prescript text
<code>postscript</code>	<code>string</code>	the postscript text

The entry `dash` is optional.

11.2.3 text

FIELD	TYPE	EXPLANATION
<code>text</code>	<code>string</code>	the text



font	string	font tfm name
dsize	number	font size
color	table	the object's color
width	number	
height	number	
depth	number	
transform	table	a text transformation
prescript	string	the prescript text
postscript	string	the postscript text

11.2.4 special

FIELD	TYPE	EXPLANATION
prescript	string	special text

11.2.5 start_bounds, start_clip

FIELD	TYPE	EXPLANATION
path	table	the list of knots

11.2.5.1 stop_bounds, stop_clip

Here are no fields available.

11.3 Subsidiary table formats

11.3.1 Paths and pens

Paths and pens (that are really just a special type of paths as far as MPLIB is concerned) are represented by an array where each entry is a table that represents a knot.

FIELD	TYPE	EXPLANATION
left_type	string	when present: endpoint, but usually absent
right_type	string	like left_type
x_coord	number	X coordinate of this knot
y_coord	number	Y coordinate of this knot
left_x	number	X coordinate of the precontrol point of this knot
left_y	number	Y coordinate of the precontrol point of this knot
right_x	number	X coordinate of the postcontrol point of this knot
right_y	number	Y coordinate of the postcontrol point of this knot

There is one special case: pens that are (possibly transformed) ellipses have an extra string-valued key type with value `elliptical` besides the array part containing the knot list.



11.3.2 Colors

A color is an integer array with 0, 1, 3 or 4 values:

FIELD	TYPE	EXPLANATION
0	marking only	no values
1	greyscale	one value in the range (0, 1), 'black' is 0
3	RGB	three values in the range (0, 1), 'black' is 0, 0, 0
4	CMYK	four values in the range (0, 1), 'black' is 0, 0, 0, 1

If the color model of the internal object was uninitialized, then it was initialized to the values representing 'black' in the colorspace `defaultcolormodel` that was in effect at the time of the shipout.

11.3.3 Transforms

Each transform is a six-item array.

INDEX	TYPE	EXPLANATION
1	number	represents x
2	number	represents y
3	number	represents xx
4	number	represents yx
5	number	represents xy
6	number	represents yy

Note that the translation (index 1 and 2) comes first. This differs from the ordering in POSTSCRIPT, where the translation comes last.

11.3.4 Dashes

Each dash is two-item hash, using the same model as POSTSCRIPT for the representation of the dashlist. `dashes` is an array of 'on' and 'off', values, and `offset` is the phase of the pattern.

FIELD	TYPE	EXPLANATION
<code>dashes</code>	hash	an array of on-off numbers
<code>offset</code>	number	the starting offset value

11.3.5 Pens and `pen_info`

There is helper function (`pen_info(obj)`) that returns a table containing a bunch of vital characteristics of the used pen (all values are floats):

FIELD	TYPE	EXPLANATION
<code>width</code>	number	width of the pen
<code>sx</code>	number	x scale



rx	number	x y multiplier
ry	number	y x multiplier
sy	number	y scale
tx	number	x offset
ty	number	y offset

11.4 Accessors

11.4.1 Character size information

These functions find the size of a glyph in a defined font. The fontname is the same name as the argument to `infont`; the char is a glyph id in the range 0 to 255; the returned w is in AFM units.

```
<number> w = char_width(mp,<string> fontname, <number> char)
<number> h = char_height(mp,<string> fontname, <number> char)
<number> d = char_depth(mp,<string> fontname, <number> char)
```

These are rather useless and might become obsolete.

11.4.2 `get_[boolean|numeric|string|path]`

When a script call brings you from the METAPost run (temporarily) back to Lua you can access variables, but only if they are known (so for instance anonymous capsules like loop variables are not accessible).

```
<boolean> w = get_boolean(mp,<string> name)
<number>   n = get_numeric(mp,<string> name)
<string>   s = get_string (mp,<string> name)
<table>    p = get_path   (mp,<string> name)
```

The path is returned as a table with subtables that have six numbers: the coordinates of the point, pre- and postcontrol. A cycle field indicates if a path is cyclic.





12 The PDF related libraries

12.1 The pdf library

12.1.1 Introduction

The pdf library replaces the epdf library and provides an interface to PDF files. It uses the same code as is used for PDF image inclusion. The pplib library by Paweł Jackowski replaces the poppler (derived from xpdf) library.

A PDF file is basically a tree of objects and one descends into the tree via dictionaries (key/value) and arrays (index/value). There are a few topmost dictionaries that start at root that are accessed more directly.

Although everything in PDF is basically an object we only wrap a few in so called userdata LUA objects.

TYPE	MAPPING
null	nil
boolean	boolean
integer	integer
float	number
name	string
string	string
array	array userdata
dictionary	dictionary userdata
stream	stream userdata (with related dictionary)
reference	reference userdata

The regular getters return these LUA data types but one can also get more detailed information.

12.1.2 open, openfile, new, getstatus, close, unencrypt

A document is loaded from a file (by name or handle) or string:

```
<pdf document> = pdf.open(filename)
<pdf document> = pdf.openfile(filehandle)
<pdf document> = pdf.new(somestring,somelength)
```

Such a document is closed with:

```
pdf.close(<pdf document>)
```

You can check if a document opened well by:

```
pdf.getstatus(<pdf document>)
```



The returned codes are:

VALUE	EXPLANATION
-2	the document failed to open
-1	the document is (still) protected
0	the document is not encrypted
2	the document has been unencrypted

An encrypted document can be unencrypted by the next command where instead of either password you can give nil:

```
pdfc.unencrypt(<pdfc document>,userpassword,ownerpassword)
```

12.1.3 getsize, getversion, getnofobjects, getnofpages

A successfully opened document can provide some information:

```
bytes = getsize(<pdfc document>)
major, minor = getversion(<pdfc document>)
n = getnofobjects(<pdfc document>)
n = getnofpages(<pdfc document>)
bytes, waste = getnofpages(<pdfc document>)
```

12.1.4 get[catalog|trailer|info]

For accessing the document structure you start with the so called catalog, a dictionary:

```
<pdfc dictionary> = pdfc.getcatalog(<pdfc document>)
```

The other two root dictionaries are accessed with:

```
<pdfc dictionary> = pdfc.gettrailer(<pdfc document>)
<pdfc dictionary> = pdfc.getinfo(<pdfc document>)
```

12.1.5 getpage, getbox

A specific page can conveniently be reached with the next command, which returns a dictionary. The first argument is to be a page dictionary.

```
<pdfc dictionary> = pdfc.getpage(<pdfc document>,pagenumber)
```

Another convenience command gives you the (bounding) box of a (normally page) which can be inherited from the document itself. An example of a valid box name is MediaBox.

```
pages = pdfc.getbox(<pdfc dictionary>,boxname)
```



12.1.6 get[string|integer|number|boolean|name]

Common values in dictionaries and arrays are strings, integers, floats, booleans and names (which are also strings) and these are also normal LUA objects:

```
s = getstring (<pdf array|dictionary>,index|key)
i = getinteger(<pdf array|dictionary>,index|key)
n = getnumber (<pdf array|dictionary>,index|key)
b = getboolean(<pdf array|dictionary>,index|key)
n = getname   (<pdf array|dictionary>,index|key)
```

12.1.7 get[from][dictionary|array|stream]

Normally you will use an index in an array and key in a dictionary but dictionaries also accept an index. The size of an array or dictionary is available with the usual # operator.

```
<pdf dictionary> = getdictionary(<pdf array|dictionary>,index|key)
<pdf array>      = getarray      (<pdf array|dictionary>,index|key)
<pdf stream>,
<pdf dictionary> = getstream     (<pdf array|dictionary>,index|key)
```

These commands return dictionaries, arrays and streams, which are dictionaries with a blob of data attached.

Before we come to an alternative access mode, we mention that the objects provide access in a different way too, for instance this is valid:

```
print(pdf.open("foo.pdf").Catalog.Type)
```

At the topmost level there are Catalog, Info, Trailer and Pages, so this is also okay:

```
print(pdf.open("foo.pdf").Pages[1])
```

12.1.8 [open|close|readfrom|whole|]stream

Streams are sort of special. When your index or key hits a stream you get back a stream object and dictionary object. The dictionary you can access in the usual way and for the stream there are the following methods:

```
okay    = openstream(<pdf stream>,[decode])
         closestream(<pdf stream>)
str, n = readfromstream(<pdf stream>)
str, n = readwholestream(<pdf stream>,[decode])
```

You either read in chunks, or you ask for the whole. When reading in chunks, you need to open and close the stream yourself. The n value indicates the length read. The decode parameter controls if the stream data gets uncompressed.



As with dictionaries, you can access fields in a stream dictionary in the usual LUA way too. You get the content when you 'call' the stream. You can pass a boolean that indicates if the stream has to be decompressed.

12.1.9 getfrom[dictionary|array]

In addition to the interface described before, there is also a bit lower level interface available.

```
key, type, value, detail = getfromdictionary(<pdf dictionary>,index)
type, value, detail = getfromarray(<pdf array>,index)
```

TYPE	MEANING	VALUE	DETAIL
0	none	nil	
1	null	nil	
2	boolean	boolean	
3	integer	integer	
4	number	float	
5	name	string	
6	string	string	hex
7	array	arrayobject	size
8	dictionary	dictionaryobject	size
9	stream	streamobject	dictionary size
10	reference	integer	

A hex string is (in the PDF file) surrounded by <> while plain strings are bounded by <>.

12.1.10 [dictionary|array]totable

All entries in a dictionary or table can be fetched with the following commands where the return values are a hashed or indexed table.

```
hash = dictionarytotable(<pdf dictionary>)
list = arraytotable(<pdf array>)
```

You can get a list of pages with:

```
{ { <pdf dictionary>, size, objnum }, ... } = pagestotable(<pdf document>)
```

12.1.11 getfromreference

Because you can have unresolved references, a reference object can be resolved with:

```
type, <pdf dictionary|array|stream>, detail = getfromreference(<pdf reference>)
```

So, as second value you get back a new pdf userdata object that you can query.



12.2 Memory streams

The `pdfc.new` that takes three arguments:

VALUE	EXPLANATION
<code>stream</code>	this is a (in low level LUA speak) light userdata object, i.e. a pointer to a sequence of bytes
<code>length</code>	this is the length of the stream in bytes (the stream can have embedded zeros)
<code>name</code>	optional, this is a unique identifier that is used for hashing the stream, so that multiple doesn't use more memory

The third argument is optional. When it is not given the function will return an `pdfc` document object as with a regular file, otherwise it will return a filename that can be used elsewhere (e.g. in the `image` library) to reference the stream as pseudo file.

Instead of a light userdata stream (which is actually fragile but handy when you come from a library) you can also pass a LUA string, in which case the given length is (at most) the string length.

The function returns an `pdfc` object and a string. The string can be used in the `img` library instead of a filename. You need to prevent garbage collection of the object when you use it as image (for instance by storing it somewhere).

Both the memory stream and it's use in the `image` library is experimental and can change. In case you wonder where this can be used: when you use the `swiglib` library for `graphicmagick`, it can return such a userdata object. This permits conversion in memory and passing the result directly to the backend. This might save some runtime in one-pass workflows. This feature is currently not meant for production and we might come up with a better implementation.

12.3 The pdfscanner library

This library is not available in `LUAMETATEX`.





13 Extra libraries

13.1 Introduction

The libraries can be grouped in categories like fonts, languages, T_EX, METAPOST, PDF, etc. There are however also some that are more general purpose and these are discussed here.

13.2 File and string readers: **fio** and **sio**

This library provides a set of functions for reading numbers from a file and in addition to the regular `io` library functions. The following work on normal LUA file handles.

NAME	ARGUMENTS	RESULTS
<code>readcardinal1</code>	<code>(f)</code>	a 1 byte unsigned integer
<code>readcardinal2</code>	<code>(f)</code>	a 2 byte unsigned integer
<code>readcardinal3</code>	<code>(f)</code>	a 3 byte unsigned integer
<code>readcardinal4</code>	<code>(f)</code>	a 4 byte unsigned integer
<code>readcardinaltable</code>	<code>(f,n,b)</code>	n cardinals of b bytes
<code>readinteger1</code>	<code>(f)</code>	a 1 byte signed integer
<code>readinteger2</code>	<code>(f)</code>	a 2 byte signed integer
<code>readinteger3</code>	<code>(f)</code>	a 3 byte signed integer
<code>readinteger4</code>	<code>(f)</code>	a 4 byte signed integer
<code>readintegertable</code>	<code>(f,n,b)</code>	n integers of b bytes
<code>readfixed2</code>	<code>(f)</code>	a 2 byte float (used in font files)
<code>readfixed4</code>	<code>(f)</code>	a 4 byte float (used in font files)
<code>read2dot14</code>	<code>(f)</code>	a 2 byte float (used in font files)
<code>setposition</code>	<code>(f,p)</code>	goto position p
<code>getposition</code>	<code>(f)</code>	get the current position
<code>skipposition</code>	<code>(f,n)</code>	skip n positions
<code>readbytes</code>	<code>(f,n)</code>	n bytes
<code>readbytetable</code>	<code>(f,n)</code>	n bytes

When relevant there are also variants that end with `le` that do it the little endian way.

A similar set of function as in the `fio` library is available in the `sio` library: `sio.readcardinal1`, `sio.readcardinal2`, `sio.readcardinal3`, `sio.readcardinal4`, `sio.readcardinaltable`, `sio.readinteger1`, `sio.readinteger2`, `sio.readinteger3`, `sio.readinteger4`, `sio.readintegertable`, `sio.readfixed2`, `sio.readfixed4`, `sio.read2dot14`, `sio.setposition`, `sio.getposition`, `sio.skipposition`, `sio.readbytes` and `sio.readbytetable`. Here the first argument is a string instead of a file handle.

13.3 md5

NAME	ARGUMENTS	RESULTS
<code>sum</code>		



hex
HEX

13.4 sha2

NAME	ARGUMENTS	RESULTS
digest256		
digest384		
digest512		

13.5 xzip

NAME	ARGUMENTS	RESULTS
compress		
decompress		
adler32		
crc32		

13.6 xmath

This library just opens up standard C math library and the main reason for it being there is that it permits advanced graphics in METAPOST (via the LUA interface). There are three constant values:

NAME	ARGUMENTS	RESULTS
inf	–	inf
nan	–	nan
pi	–	3.1415926535898

and a lot of functions:

NAME	ARGUMENTS	RESULTS
acos	(a)	
acosh	(a)	
asin	(a)	
asinh	(a)	
atan	(a[,b])	
atan2	(a[,b])	
atanh	(a)	
cbrt	(a)	
ceil	(a)	
copysign	(a,b)	
cos	(a)	
cosh	(a)	



deg	(a)
erf	(a)
erfc	(a)
exp	(a)
exp2	(a)
expm1	(a)
fabs	(a)
fdim	(a,b)
floor	(a)
fma	(a,b,c)
fmax	(...)
fmin	(...)
fmod	(a,b)
frexp	(a,b)
gamma	(a)
hypot	(a,b)
isfinite	(a)
isinf	(a)
isnan	(a)
isnormal	(a)
j0	(a)
j1	(a)
jn	(a,b)
ldexp	(a,b)
lgamma	(a)
l0	(a)
l1	(a)
ln	(a,b)
log	(a[,b])
log10	(a)
log1p	(a)
log2	(a)
logb	(a)
modf	(a,b)
nearbyint	(a)
nextafter	(a,b)
pow	(a,b)
rad	(a)
remainder	(a,b)
remquo	(a,b)
round	(a)
scalbn	(a,b)
sin	(a)
sinh	(a)
sqrt	(a)
tan	(a)



tanh	(a)
tgamma	(a)
trunc	(a)
y0	(a)
y1	(a)
yn	(a)

13.7 xcomplex

LUAMETATEX also provides a complex library xcomplex. The complex number is a userdata:

NAME	ARGUMENTS	RESULTS
new	(r,i)	a complex userdata type
tostring	(z)	a string representation
topair	(z)	two numbers

There is a bunch of functions that take a complex number:

NAME	ARGUMENTS	RESULTS
abs	(a)	
arg	(a)	
imag	(a)	
real	(a)	
onj	(a)	
proj	(a)	
exp"	(a)	
log	(a)	
sqrt	(a)	
pow	(a,b)	
sin	(a)	
cos	(a)	
tan	(a)	
asin	(a)	
acos	(a)	
atan	(a)	
sinh	(a)	
cosh	(a)	
tanh	(a)	
asinh	(a)	
acosh	(a)	
atanh	(a)	

These are accompanied by libcerf functions:

NAME	ARGUMENTS	RESULTS
erf	(a)	The complex error function erf(z)



erfc	(a)	The complex complementary error function $\text{erfc}(z) = 1 - \text{erf}(z)$
erfcx	(a)	The underflow-compensating function $\text{erfcx}(z) = \exp(z^2) \text{erfc}(z)$
erfi	(a)	The imaginary error function $\text{erfi}(z) = -i \text{erf}(iz)$
dawson	(a)	Dawson's integral $D(z) = \sqrt{\pi}/2 * \exp(-z^2) * \text{erfi}(z)$
voigt	(a,b,c)	The convolution of a Gaussian and a Lorentzian
voigt_hwhm	(a,b)	The half width at half maximum of the Voigt profile

13.8 lfs

The original lfs module has been adapted a bit to our needs but for practical reasons we kept the namespace. This module will probably evolve a bit over time.

NAME	ARGUMENTS	RESULTS
attributes	(name)	
chdir	(name)	
currentdir	()	
dir	(name)	name, mode, size and mtime
mkdir	(name)	
rmdir	(name)	
touch	(name)	
link	(name)	
symlinkattributes	(name)	
isdir	(name)	
isfile	(name)	
iswriteabledir	(name)	
iswriteablefile	(name)	
isreadabledir	(name)	
isreadablefile	(name)	

The dir function is a traverser which in addition to the name returns some more properties. Keep in mind that the traverser loops over a directory and that it doesn't run well when used nested. This is a side effect of the operating system. It is also the reason why we return some properties because querying them via attributes would interfere badly.

The following attributes are returned by attributes:

NAME	VALUE
mode	
size	
modification	
access	
change	
permissions	
nlink	



13.9 pngdecode

This module is experimental and used in image inclusion. It is not some general purpose module and is supposed to be used in a very controlled way. The interfaces might evolve.

NAME	ARGUMENTS	RESULTS
applyfilter	(str,nx,ny,slice)	string
splitmask	(str,nx,ny,bpp,bytes)	string
interlace	(str,nx,ny,slice,pass)	string
expand	(str,nx,ny,parts,xline,factor)	string

13.10 basexx

Some more experimental helpers:

NAME	ARGUMENTS	RESULTS
encode16	(str[,newline])	string
decode16	(str)	string
encode64	(str[,newline])	string
decode64	(str)	string
encode85	(str[,newline])	string
decode85	(str)	string
encodeRL	(str)	string
decodeRL	(str)	string
encodeLZW	(str[,defaults])	string
decodeLZW	(str[,defaults])	string

13.11 Multibyte string functions

The string library has a few extra functions, for example `string.explode`. This function takes upto two arguments: `string.explode(s[,m])` and returns an array containing the string argument `s` split into sub-strings based on the value of the string argument `m`. The second argument is a string that is either empty (this splits the string into characters), a single character (this splits on each occurrence of that character, possibly introducing empty strings), or a single character followed by the plus sign `+` (this special version does not create empty sub-strings). The default value for `m` is `' '` (multiple spaces). Note: `m` is not hidden by surrounding braces as it would be if this function was written in \TeX macros.

The string library also has six extra iterators that return strings piecemeal: `string.utfvalues`, `string.utfcharacters`, `string.characters`, `string.characterpairs`, `string.bytes` and `string.bytepairs`.

- ▶ `string.utfvalues(s)`: an integer value in the UNICODE range
- ▶ `string.utfcharacters(s)`: a string with a single UTF-8 token in it
- ▶ `string.cWharacters(s)`: a string containing one byte



- ▶ `string.characterpairs(s)`: two strings each containing one byte or an empty second string if the string length was odd
- ▶ `string.bytes(s)`: a single byte value
- ▶ `string.bytepairs(s)`: two byte values or nil instead of a number as its second return value if the string length was odd

The `string.characterpairs()` and `string.bytepairs()` iterators are useful especially in the conversion of UTF16 encoded data into UTF8.

There is also a two-argument form of `string.dump()`. The second argument is a boolean which, if true, strips the symbols from the dumped data. This matches an extension made in `lua.jit`. This is typically a function that gets adapted as LUA itself progresses.

The `string` library functions `len`, `lower`, `sub` etc. are not UNICODE-aware. For strings in the UTF8 encoding, i.e., strings containing characters above code point 127, the corresponding functions from the `slnunicode` library can be used, e.g., `unicode.utf8.len`, `unicode.utf8.lower` etc. The exceptions are `unicode.utf8.find`, that always returns byte positions in a string, and `unicode.utf8.match` and `unicode.utf8.gmatch`. While the latter two functions in general *are* UNICODE-aware, they fall-back to non-UNICODE-aware behavior when using the empty capture `()` but other captures work as expected. For the interpretation of character classes in `unicode.utf8` functions refer to the library sources at <http://luaforge.net/projects/sln>.

Version 5.3 of LUA provides some native UTF8 support but we have added a few similar helpers too: `string.utfvalue`, `string.utfcharacter` and `string.utflength`.

- ▶ `string.utfvalue(s)`: returns the codepoints of the characters in the given string
- ▶ `string.utfcharacter(c, ...)`: returns a string with the characters of the given code points
- ▶ `string.utflength(s)`: returns the length of the given string

These three functions are relative fast and don't do much checking. They can be used as building blocks for other helpers.

13.12 Extra os library functions

The `os` library has a few extra functions and variables: `os.selfdir`, `os.selfarg`, `os.setenv`, `os.env`, `os.gettimeofday`, `os.type`, `os.name` and `os.uname`, that we will discuss here. There are also some time related helpers in the `lua` namespace.

- ▶ `os.selfdir` is a variable that holds the directory path of the actual executable. For example: `\directlua{tex.sprint(os.selfdir)}`.
- ▶ `os.selfarg` is a table with the command line arguments.
- ▶ `os.setenv(key,value)` sets a variable in the environment. Passing nil instead of a value string will remove the variable.
- ▶ `os.env` is a hash table containing a dump of the variables and values in the process environment at the start of the run. It is writeable, but the actual environment is *not* updated automatically.
- ▶ `os.gettimeofday` returns the current 'UNIX time', but as a float. Keep in mind that there might be platforms where this function is available.
- ▶ `os.type` is a string that gives a global indication of the class of operating system. The possible values are currently `windows`, `unix`, and `msdos` (you are unlikely to find this value 'in the wild').



- ▶ `os.name` is a string that gives a more precise indication of the operating system. These possible values are not yet fixed, and for `os.type` values `windows` and `msdos`, the `os.name` values are simply `windows` and `msdos`

The list for the type `unix` is more precise: `linux`, `freebsd`, `kfreebsd`, `cygwin`, `openbsd`, `solaris`, `sunos` (pre-solaris), `hpux`, `irix`, `macosx`, `gnu` (hurd), `bsd` (unknown, but BSD-like), `sysv`, `generic` (unknown). But ... we only provide LUAMETATEX binaries for the mainstream variants.

Officially we only support mainstream systems: MS WINDOWS, LINUX, FREEBSD and OS-X. Of course one can build LUAMETATEX for other systems, in which case one has to check the above.

- ▶ `os.uname` returns a table with specific operating system information acquired at runtime. The keys in the returned table are all string values, and their names are: `sysname`, `machine`, `release`, `version`, and `nodename`.

13.13 The lua library functions

The lua library provide some general helpers.

- ▶ The `newtable` and `newindex` functions can be used to create tables with space reserved beforehand for the given amount of entries.
- ▶ The `getstacktop` function returns a number that can be used for diagnostic purposes.
- ▶ The functions `getruntime`, `getcurrenttime`, `getpreciseticks` and `getpreciseseconds` return what their name suggests.
- ▶ On MS WINDOWS the `getcodepage` function returns two numbers, one for the command handler and one for the graphical user interface.
- ▶ The name of the startup file is reported by `getstartupfile`.
- ▶ The LUA version is reported by `getversion`.
- ▶ The `lua.openfile` function can be used instead of `io.open`. On MS WINDOWS it will convert the filename to a so called wide one which means that filenames in UTF8 encoding will work ok. On the other hand, names given in the codepage won't.



14 Primitives aka commands

14.1 Introduction

The starting point of L^AT_EX is P_DF_TE_X, which itself contains regular T_EX and ϵ -T_EX. Because directional support was needed we also took some code from ALEPH (OMEGA). In a later stage the backend specific commands were isolated in its own namespace which resulted in a cleaner code base where the backend code no longer was interwoven with the normal frontend primitives. We also promoted some generic constructs (like box resources and directions) to core functionality.

Some of the P_DF_TE_X support primitives have been around from the start but when L^A integration became better and when a token scanner library was added, not all of those made sense as primitives. In previous chapters we already mentioned what is gone from the core. Deep down some more has changes but not all is reflected at the primitive level. Because there is still a considerable amount of new primitives, a summary is given below.

14.2 Languages

automatichyphenmode	integer
automatichyphenpenalty	integer
hyphenpenaltymode	integer
compoundhyphenmode	integer
exceptionpenalty	integer
explicithyphenpenalty	integer
hyphenationbounds	integer
hocode	charactercode
hyphenationmin	charactercode
postexhyphenchar	charactercode
posthyphenchar	charactercode
preexhyphenchar	charactercode
prehyphenchar	charactercode

14.3 Fonts

tracingfonts	integer
suppressfontnotfounderror	integer
setfontid	integer
fontid	font
efcode	font charactercode
lpcode	font charactercode
rpcode	font charactercode



14.4 Math

<code>matholdmode</code>	integer
<code>mathstyle</code>	integer
<code>matheqnogapstep</code>	integer
<code>Uskewed</code>	
<code>Uskewedwithdelims</code>	
<code>Ustartdisplaymath</code>	
<code>Ustartmath</code>	
<code>Ustopdisplaymath</code>	
<code>Ustopmath</code>	
<code>crampeddisplaystyle</code>	
<code>crampedtextstyle</code>	
<code>crampedscriptstyle</code>	
<code>crampedscriptscriptstyle</code>	
<code>Umathchardef</code>	
<code>Umathcharnumdef</code>	
<code>mathdisplayskipmode</code>	integer
<code>mathscriptsmode</code>	integer
<code>mathnolimitsmode</code>	integer
<code>mathitalicsmode</code>	integer
<code>mathrulesmode</code>	integer
<code>mathrulesfam</code>	integer
<code>mathdelimitersmode</code>	integer
<code>mathflattenmode</code>	integer
<code>mathpenaltiesmode</code>	integer
<code>mathrulethicknessmode</code>	integer
<code>mathscriptboxmode</code>	integer
<code>mathscriptcharmode</code>	integer
<code>mathsurroundmode</code>	integer
<code>nokerns</code>	integer
<code>noligs</code>	integer
<code>prebinoppenalty</code>	integer
<code>predisplaygapfactor</code>	integer
<code>prerelpenalty</code>	integer
<code>Usuperscript</code>	command
<code>Usubscript</code>	command
<code>Unosuperscript</code>	command
<code>Unosubscript</code>	command
<code>Umathcode</code>	
<code>Umathcodenum</code>	
<code>Udelcode</code>	
<code>Udelcodenum</code>	
<code>Umathaxis</code>	family dimension
<code>Umathbinbinspacing</code>	family dimension
<code>Umathbinclonespacing</code>	family dimension



<code>Umathbininnerspacing</code>	family dimension
<code>Umathbinopenspacing</code>	family dimension
<code>Umathbinopspacing</code>	family dimension
<code>Umathbinordspacing</code>	family dimension
<code>Umathbinpunctspacing</code>	family dimension
<code>Umathbinrelspacing</code>	family dimension
<code>Umathclosebinspacing</code>	family dimension
<code>Umathcloseclosespacing</code>	family dimension
<code>Umathcloseinnerspacing</code>	family dimension
<code>Umathcloseopenspacing</code>	family dimension
<code>Umathcloseopspacing</code>	family dimension
<code>Umathcloseordspacing</code>	family dimension
<code>Umathclosepunctspacing</code>	family dimension
<code>Umathcloserelspacing</code>	family dimension
<code>Umathconnectoroverlapmin</code>	family dimension
<code>Umathfractiondelsize</code>	family dimension
<code>Umathfractiondenomdown</code>	family dimension
<code>Umathfractiondenomvgap</code>	family dimension
<code>Umathfractionnumup</code>	family dimension
<code>Umathfractionnumvgap</code>	family dimension
<code>Umathfractionrule</code>	family dimension
<code>Umathinnerbinspacing</code>	family dimension
<code>Umathinnerclosespacing</code>	family dimension
<code>Umathinnerinnerspacing</code>	family dimension
<code>Umathinneropenspacing</code>	family dimension
<code>Umathinneropspacing</code>	family dimension
<code>Umathinnerordspacing</code>	family dimension
<code>Umathinnerpunctspacing</code>	family dimension
<code>Umathinnerrelspacing</code>	family dimension
<code>Umathlimitabovebgap</code>	family dimension
<code>Umathlimitabovekern</code>	family dimension
<code>Umathlimitabovevgap</code>	family dimension
<code>Umathlimitbelowbgap</code>	family dimension
<code>Umathlimitbelowkern</code>	family dimension
<code>Umathlimitbelowvgap</code>	family dimension
<code>Umathnolimitsubfactor</code>	family dimension
<code>Umathnolimitsupfactor</code>	family dimension
<code>Umathopbinspacing</code>	family dimension
<code>Umathopclosespacing</code>	family dimension
<code>Umathopenbinspacing</code>	family dimension
<code>Umathopenclosespacing</code>	family dimension
<code>Umathopeninnerspacing</code>	family dimension
<code>Umathopenopenspacing</code>	family dimension
<code>Umathopenopspacing</code>	family dimension
<code>Umathopenordspacing</code>	family dimension
<code>Umathopenpunctspacing</code>	family dimension



Umathopenrelspacing	family dimension
Umathoperatorsize	family dimension
Umathopinnerspacing	family dimension
Umathopopenspacing	family dimension
Umathopopspacing	family dimension
Umathopordspacing	family dimension
Umathoppunctspacing	family dimension
Umathoprelspacing	family dimension
Umathordbinspacing	family dimension
Umathordclosespacing	family dimension
Umathordinnerspacing	family dimension
Umathordopenspacing	family dimension
Umathordopspacing	family dimension
Umathordordspacing	family dimension
Umathordpunctspacing	family dimension
Umathordrelspacing	family dimension
Umathoverbarkern	family dimension
Umathoverbarrule	family dimension
Umathoverbarvgap	family dimension
Umathoverdelimiterbgap	family dimension
Umathoverdelimitervgap	family dimension
Umathpunctbinspacing	family dimension
Umathpunctclosespacing	family dimension
Umathpunctinnerspacing	family dimension
Umathpunctopenspacing	family dimension
Umathpunctopspacing	family dimension
Umathpunctordspacing	family dimension
Umathpunctpunctspacing	family dimension
Umathpunctrelspacing	family dimension
Umathquad	family dimension
Umathradicaldegreeafter	family dimension
Umathradicaldegreebefore	family dimension
Umathradicaldegreeraise	family dimension
Umathradicalkern	family dimension
Umathradicalrule	family dimension
Umathradicalvgap	family dimension
Umathrelbinspacing	family dimension
Umathrelclosespacing	family dimension
Umathrelinnerspacing	family dimension
Umathrelopenspacing	family dimension
Umathrelopspacing	family dimension
Umathrelordspacing	family dimension
Umathrelpunctspacing	family dimension
Umathrelrelspacing	family dimension
Umathskewedfractionhgap	family dimension
Umathskewedfractionvgap	family dimension



<code>Umathspaceafterscript</code>	family dimension
<code>Umathstackdenomdown</code>	family dimension
<code>Umathstacknumup</code>	family dimension
<code>Umathstackvgap</code>	family dimension
<code>Umathsubshiftdown</code>	family dimension
<code>Umathsubshiftdrop</code>	family dimension
<code>Umathsubsupshiftdown</code>	family dimension
<code>Umathsubsupvgap</code>	family dimension
<code>Umathsubtopmax</code>	family dimension
<code>Umathsupbottommin</code>	family dimension
<code>Umathsupshiftdrop</code>	family dimension
<code>Umathsupshiftdown</code>	family dimension
<code>Umathsupsubbottommax</code>	family dimension
<code>Umathunderbarkern</code>	family dimension
<code>Umathunderbarrule</code>	family dimension
<code>Umathunderbarvgap</code>	family dimension
<code>Umathunderdelimterbgap</code>	family dimension
<code>Umathunderdelimtervgap</code>	family dimension
<code>Udelimiter</code>	
<code>Umathaccent</code>	
<code>Umathchar</code>	
<code>Umathcharnum</code>	
<code>Ustack</code>	
<code>Uradical</code>	
<code>Uroot</code>	
<code>Uunderdelimiter</code>	
<code>Uoverdelimiter</code>	
<code>Udelimiterunder</code>	
<code>Udelimiterover</code>	
<code>Uhextensible</code>	
<code>Uchar</code>	
<code>Umathcharclass</code>	
<code>Umathcharfam</code>	
<code>Umathcharslot</code>	
<code>Uleft</code>	
<code>Umiddle</code>	
<code>Uright</code>	
<code>Uvextensible</code>	

14.5 Boxes and directions

<code>pardirection</code>	direction
<code>textdirection</code>	direction
<code>mathdirection</code>	direction
<code>linedirection</code>	direction
<code>breakafterdirmode</code>	integer



shapemode	integer	
fixupboxesmode	integer	
boxdirection	box direction	
boxorientation	box orientation	rotation over 90, 180, 270 degrees
boxxoffset	box xoffset	leaves dimensions untouched
boxyoffset	box yoffset	leaves dimensions untouched
boxxmove	box xmove	offsets that reflect on dimensions
boxymove	box ymove	offsets that reflect on dimensions
boxtotal	box ht+dp	height plus depth (and when assigned halves)
boxattr	box attr value	(sets) a specific attribute to a value

14.6 Scanning

aftergrouped	text	like aftergroup but for given list
alignmark		equivalent to hash token
aligntab		equivalent to tab token
begincsname	command	variant of \csname that ignores unde- fined commands
catcodetable	integer	switch to catcode table
csstring	command	the command without preceding es- cape character
endlocalcontrol	command	switches back to the main control loop
etoksapp	tokenregister text	append expanded text to given token- register
etokspre	tokenregister text	prepend expanded text to given token- register
expanded	text	expands the given text
frozen		
futureexpand	command command	expands second or third token depend- ing on first match
futureexpandis	command command	as futureexpand but also skips pars
futureexpandisap	command command	same as idem but doesn't push back skipped spaces
gtoksapp	tokenregister text	globally append text to given token- register
gtokspre	tokenregister text	globally prepend text to given token- register
ifabsdim	dimension <=> dimension	test the absolute value of the given di- mension
ifabsnum	integer <=> integer	test the absolute value of the given in- teger
ifcondition	command	assume the next token is a test (so skip as if)
ifdimen	possibly a dimension	acts like an \ifcase with 1 for valid and 2 for invalid



<code>ifincsname</code>	command	check if we're inside a csname expansion
<code>ifnumval</code>		
<code>ifdimval</code>		
<code>ifchknum</code>		
<code>ifchkdim</code>		
<code>ifcmpnum</code>		
<code>ifcmpdim</code>		
<code>ifusercmd</code>	command	
<code>ifprotected</code>	command	
<code>iffrozen</code>	command	
<code>iftok</code>		
<code>ifcstok</code>		
<code>internalcodesmode</code>	integer	
<code>immediateassigned</code>	command	(todo) expand the following assignment now
<code>immediateassignment</code>	command	(todo) expand the following assignment now
<code>initcatcodetable</code>	integer	initialize catcode table
<code>lastnamedcs</code>	command	last found command of <code>\ifcsname</code> construction
<code>nospaces</code>	integer	don't inject spaces
<code>orelse</code>	condition	
<code>pxdimen</code>	dimension	multiplier for the px unit
<code>savecatcodetable</code>	integer	save catcode table
<code>scantextokens</code>	text	<code>\scantokens</code> without file side effects
<code>suppressifcsnameerror</code>	integer	recover from issues in csname testing
<code>suppresslongerror</code>	integer	make <code>\long</code> a nop
<code>suppressmathparerror</code>	integer	accepts <code>\par</code> and empty lines in math
<code>suppressoutererror</code>	integer	make <code>\outer</code> a nop
<code>suppressprimitiveerror</code>	integer	don't report an invalid <code>\primitive</code>
<code>toksapp</code>	tokenregister text	append text to given tokenregister
<code>tokspre</code>	tokenregister text	prepend text to given tokenregister
<code>xtoksapp</code>	tokenregister text	globally append expanded text to given tokenregister
<code>xtokspre</code>	tokenregister text	globally prepend expanded text to given tokenregister
<code>letfrozen</code>	macro	sets the frozen property of a macro
<code>letprotected</code>	macro	sets the protected property of a macro
<code>unletfrozen</code>	macro	unsets the frozen property of a macro
<code>unletprotected</code>	macro	unsets the protected property of a macro

14.7 Typesetting

<code>protrudechars</code>	integer
<code>localbrokenpenalty</code>	integer



localinterlinepenalty	integer
adjustspacing	integer
boundary	command
noboundary	command
protrusionboundary	command
wordboundary	command
leftghost	charactercode
rightghost	charactercode
nohrule	command
novrule	command
insertht	number
quitvmode	command
leftmarginkern	dimension
rightmarginkern	dimension
localleftbox	box
localrightbox	box
gleaders	command

14.8 LUA

luacopyinputnodes integer
 luadef
 luabytecodecall
 luafunctioncall
 latelua
 lateluafunction
 luabytecode
 luaescapestring
 luafunction

14.9 Management

outputbox integer
 clearmarks
 attribute
 glet
 letcharcode
 attributedef

14.10 Miscellaneous

luatexversion
 formatname
 luatexbanner
 luatexrevision



Topics

a

ALEPH 21, 50
adjust 93
attributes 32, 33, 110, 146

b

banner 30
boundaries 62
boundary 97
boxes 11, 33, 149
 split 150
bytecodes 139

c

callbacks 127
 building pages 130
 closing files 129
 contributions 130, 132
 data files 128
 dump 136
 errors 136, 137
 files 137
 fonts 137
 format file 128
 hyphenation 134
 inserts 131
 job run 136
 jobname 129
 kerning 135
 ligature building 135
 linebreaks 131, 132, 133
 math 135
 opening files 128
 output 134
 packing 133, 134
 reader 129
 rules 134
 wrapping up 137
catcodes 37
characters
 codes 148
command line 25

conditions 44
 dimensions 42
 numbers 42
 tokens 43
configuration 165
convert commands 146
csnames 23

d

dimensions 42
direct nodes 117
directions 50, 98
discretionaries 93

e

ε -TEX 20
engines 19
errors 156
escaping 35
expansion 39

f

files
 binary 23
 names 49
 writing 49
fonts 50, 60
 current 63
 define 63
 defining 158
 extend 63
 id 63
 used 223
format 30

g

glue 94
gluespec 95
glyphs 96

h

hash 158



helpers 155
history 19
hyphenation 48

i

IO 166
images
 METAPOST 173
 mplib 173
initialization 25, 158
insertions 92

k

kerns 95
 suppress 61

l

LUA 11
 extensions 26
 interpreter 25
 libraries 26, 27
 modules 27
languages 48
last items 146
leaders 48
libraries
 lua 139
 status 140
 tex 142
 texconfig 165
 texio 166
 token 166
ligatures
 suppress 61
linebreaks 163
lists 91, 151

m

METAPOST 173
 mplib 173
macros 170
marks 39, 92
math 50, 65
 accents 81, 85
 codes 85
 cramped 68

delimiters 83, 84
extensibles 83
fences 81
flattening 87
fractions 84
italics 78
kerning 78
last line 86
limits 78
nodes 94, 98
parameters 69, 71, 150
penalties 80
radicals 82
scripts 78, 83, 86
spacing 68, 75, 76, 77
stacks 68
styles 66, 68, 86
text 87
tracing 88
 UNICODE 65
memory 22

n

nesting 152, 164
newline 23
nodes 11, 32, 89
 adjust 93
 attributes 111
 boundary 97
 direct 117
 direction 98
 discretionaries 93
 functions 105
 glue 94, 95
 glyph 96
 insertions 92
 kerns 95
 lists 91
 marks 92
 math 94, 98
 paragraphs 97, 98
 penalty 95
 properties 122
 rules 91
 text 90
numbers 42



o

OMEGA 50

output 46

p

PDF

analyze 181

memory streams 185

objects 181

pdfc 181

PDFTEX 21

pages 150, 164

paragraphs 97, 98

reset 162

parameters

internal 142

math 150

penalty 95

primitives 158

printing 153

properties 122

protrusion 62, 98

r

registers 146, 149

bytecodes 139

rules 48, 91

s

shipout 164

space 23

spaces

suppress 61

splitting 46

startupfile 139

synctex 164

t

TEX 19

tables 139

testing 27

text

math 87

tokens 43, 166

scanning 38

tracing 49

u

UNICODE 30, 31

math 65

v

version 30, 139

w

WEB2C 22





Primitives

This register contains the primitives that are mentioned in the manual. There are of course many more primitives. The L^AT_EX primitives are typeset in bold. The primitives from PDF_TE_X are not supported that way but mentioned anyway.

<code>\abovedisplayskip</code>	77	<code>\delimiter</code>	65
<code>\abovewithdelims</code>	84	<code>\detokenize</code>	168
<code>\accent</code>	41	<code>\dimen</code>	27, 31, 146
<code>\adjustspacing</code>	21, 57	<code>\dimendef</code>	31, 146, 147
<code>\adjustspacingshrink</code>	21	<code>\directlua</code>	11
<code>\adjustspacingstep</code>	21	<code>\directlua</code>	29, 34, 35, 140, 153, 158, 159
<code>\adjustspacingstretch</code>	21	<code>\discretionary</code>	12, 93
<code>\aftergrouped</code>	41	<code>\displaystyle</code>	75
<code>\alignmark</code>	39	<code>\displaywidowpenalties</code>	163
<code>\aligntab</code>	39	<code>\dp</code>	31
<code>\atop</code>	68, 70	<code>\edef</code>	36, 39, 168
<code>\atopwithdelims</code>	68	<code>\efcode</code>	21, 31, 57
<code>\attribute</code>	146	<code>\endgroup</code>	68
<code>\attributedef</code>	146, 147	<code>\endinput</code>	166
<code>\batchmode</code>	166	<code>\endlinechar</code>	20, 38, 153, 155
<code>\begincsname</code>	38	<code>\errhelp</code>	156
<code>\begingroup</code>	68	<code>\errmessage</code>	156
<code>\belowdisplayskip</code>	77	<code>\etoksapp</code>	38
<code>\boundary</code>	48, 97	<code>\etokspre</code>	38
<code>\box</code>	31	<code>\everyeof</code>	38
<code>\breakafterdirmode</code>	52	<code>\everyjob</code>	26
<code>\catcode</code>	22, 29, 31, 148	<code>\exhyphenpenalty</code>	93
<code>\catcodetable</code>	37, 153	<code>\expandafter</code>	39
<code>\char</code>	12, 31, 96	<code>\expanded</code>	21, 39
<code>\chardef</code>	31, 170, 171	<code>\fontid</code>	61
<code>\clearmarks</code>	39	<code>\formatname</code>	30, 159
<code>\clubpenalties</code>	163	<code>\frozen</code>	46
<code>\copy</code>	31	<code>\futureexpand</code>	41
<code>\count</code>	27, 31, 32, 146	<code>\futureexpandis</code>	41
<code>\countdef</code>	31, 146, 147	<code>\futureexpandisap</code>	41
<code>\crampedscriptstyle</code>	69	<code>\gleaders</code>	48
<code>\csname</code>	38	<code>\glet</code>	39
<code>\csstring</code>	38	<code>\global</code>	22
<code>\currentiftype</code>	49	<code>\gtoksapp</code>	38
<code>\currentiftype</code>	89		
<code>\delcode</code>	22, 65, 148, 149		



\gtokspre 38	\kern 12, 95
\halign 132	\lastnamedcs 38, 39
\hangindent 52	\lastnodesubtype 49
\hbox 12, 32, 46, 78, 131, 132, 133, 149	\lastnodetype 49
\hjcode 22, 31	\lastnodetype 89
\hpack 46	\lastsavedboxresourceindex 47
\hpack 48	\lastsavedimageresourceindex 47
\hrule 12	\lastsavedimageresourcepages 47
\hskip 12, 94	\latelua 140
\ht 31	\lccode 22, 31, 148
\hyphenationmin 48	\leaders 48
\hyphenchar 55	\left 81
\hyphenpenalty 93	\lefthyphenmin 48
	\leftmarginkern 21
\if 39	\letcharcode 39
\ifabsdim 21, 42	\linedir 52
\ifabsnum 21, 42	\localbrokenpenalty 97
\ifchkdim 42	\localinterlinepenalty 97
\ifchknum 42	\localleftbox 97, 132
\ifcmpdim 42	\localrightbox 97, 132
\ifcmpnum 42	\lpcode 21, 31, 57
\ifcondition 44	\luabytecode 36
\ifcstok 43	\luabytecodecall 36
\ifdimval 42	\luacopyinputnodes 154
\iffrozen 46	\luadef 35, 171
\ifincsname 21	\luaescapestring 35
\ifnumval 42	\luafunction 35, 36
\ifprotected 46	\luafunctioncall 35, 36
\iftok 43	\luatexbanner 30
\ifusercmd 46	\luatexrevision 30
\ignorepars 41	\luatexversion 30
\ignorespaces 41	
\immediateassigned 39	\mag 20
\immediateassignment 39	\mark 92
\initcatcodetable 37	\marks 31, 115
\input 128	\mathaccent 65
\insert 31, 92	\mathchar 65, 87
\interlinepenalties 163	\mathchardef 65, 87
\internalcodesmode 49	\mathchoice 67
\internalcodesmode 89	\mathcode 22, 65, 148
	\mathdelimitersmode 81
\jobname 25, 26, 30, 129	\mathdir 152
	\mathdisplayskipmode 77



<code>\matheqnogapstep</code>	80	<code>\pxdimen</code>	21
<code>\mathflattenmode</code>	87	<code>\quitvmode</code>	21
<code>\mathitalicsmode</code>	78, 81	<code>\radical</code>	65
<code>\mathnolimitsmode</code>	78	<code>\read</code>	128
<code>\mathpenaltiesmode</code>	80	<code>\relax</code>	154, 159, 169
<code>\mathscriptboxmode</code>	78, 79	<code>\right</code>	81
<code>\mathscriptcharmode</code>	79	<code>\righthyphenmin</code>	48
<code>\mathscriptsmode</code>	79, 80	<code>\rightmarginkern</code>	21
<code>\mathstyle</code>	66, 67, 68, 153	<code>\romannumeral</code>	67, 155
<code>\mathsurround</code>	75, 94	<code>\rptcode</code>	21, 31, 57
<code>\mathsurroundmode</code>	75	<code>\rule</code>	91
<code>] 75</code>		<code>\saveboxresource</code>	47
<code>\mathsurroundskip</code>	75	<code>\savecatcodetable</code>	38
<code>\maxdepth</code>	133	<code>\saveimageresource</code>	47
<code>\medmuskip</code>	77	<code>\scantextokens</code>	38
<code>\middle</code>	152	<code>\scantokens</code>	35, 38
<code>\muskip</code>	31, 76, 77, 146	<code>\scriptfont</code>	71
<code>\muskipdef</code>	31	<code>\scriptscriptfont</code>	71
<code>\newlinechar</code>	20	<code>\scriptscriptstyle</code>	82
<code>\noboundary</code>	48, 97	<code>\scriptspace</code>	74
<code>\noexpand</code>	39	<code>\scriptstyle</code>	69
<code>\nohrule</code>	48	<code>\setbox</code>	31
<code>\nokerns</code>	61	<code>\setfontid</code>	61
<code>\noligs</code>	61	<code>\sfcode</code>	22, 31, 148
<code>\nospaces</code>	62	<code>\shapemode</code>	52
<code>\novrule</code>	48	<code>\skewchar</code>	55, 82
<code>\number</code>	61, 155	<code>\skip</code>	31, 146, 147
<code>\openin</code>	128	<code>\skipdef</code>	31, 146, 147
<code>\openout</code>	49, 128	<code>\spaceskip</code>	61
<code>\orelse</code>	45	<code>\special</code>	59
<code>\output</code>	134, 142	<code>\string</code>	38
<code>\outputbox</code>	46	<code>\textdir</code>	98
<code>\over</code>	68, 70, 152	<code>\textdir</code>	152
<code>\overline</code>	69	<code>\textdirection</code>	12
<code>\overwithdelims</code>	68	<code>\textdirection</code>	51
<code>\par</code>	33, 41, 131	<code>\textfont</code>	71, 87
<code>\parfillskip</code>	131, 163	<code>\textstyle</code>	67
<code>\parindent</code>	142	<code>\the</code>	30, 32, 61, 142, 145, 147, 153
<code>\parshape</code>	52	<code>\thickmuskip</code>	77
<code>\penalty</code>	95	<code>\thinmuskip</code>	77
<code>\predisplaygapfactor</code>	86	<code>\toks</code>	31, 145, 146, 147, 153
<code>\protrudechars</code>	21, 57	<code>\toksapp</code>	38
<code>\protrusionboundary</code>	48, 97		



<code>\toksdef</code>	31, 146, 147	<code>\Umathfractiondelsize</code>	70
<code>\tokspre</code>	38	<code>\Umathfractiondenomdown</code>	70
<code>\tpack</code>	46	<code>\Umathfractiondenomvgap</code>	70
<code>\tpack</code>	48	<code>\Umathfractionnumup</code>	70
<code>\tracingassigns</code>	20, 22	<code>\Umathfractionnumvgap</code>	70
<code>\tracingcommands</code>	142	<code>\Umathfractionrule</code>	70
<code>\tracingfonts</code>	21, 49	<code>\Umathinnerbinspacing</code>	76
<code>\tracingnesting</code>	157	<code>\Umathinnerclosespacing</code>	76
<code>\tracingonline</code>	49	<code>\Umathinnerinnerspacing</code>	76
<code>\tracingrestores</code>	20, 22	<code>\Umathinneropenspacing</code>	76
<code>\Uchar</code>	31	<code>\Umathinneropspacing</code>	76
<code>\Udelcode</code>	66, 149	<code>\Umathinnerordspacing</code>	76
<code>\Udelcodenum</code>	66	<code>\Umathinnerpunctspacing</code>	76
<code>\Udelimiter</code>	66	<code>\Umathinnerrelspacing</code>	76
<code>\Udelimiterover</code>	66, 83	<code>\Umathlimitabovebgap</code>	70
<code>\Udelimiterunder</code>	66, 83	<code>\Umathlimitabovekern</code>	70, 74
<code>\Uhexensible</code>	83	<code>\Umathlimitabovevgap</code>	70
<code>\Umathaccent</code>	66, 81	<code>\Umathlimitbelowbgap</code>	70
<code>\Umathaxis</code>	70	<code>\Umathlimitbelowkern</code>	70, 74
<code>\Umathbinbinspacing</code>	76	<code>\Umathlimitbelowvgap</code>	70
<code>\Umathbinclosespacing</code>	76	<code>\Umathnolimitsubfactor</code>	78
<code>\Umathbininnerspacing</code>	76	<code>\Umathnolimitsupfactor</code>	78
<code>\Umathbinopspacing</code>	76	<code>\Umathopbinspacing</code>	76
<code>\Umathbinopspacing</code>	76	<code>\Umathopclosepacing</code>	76
<code>\Umathbinordspacing</code>	76	<code>\Umathopenbinspacing</code>	76
<code>\Umathbinpunctspacing</code>	76	<code>\Umathopenclosepacing</code>	76
<code>\Umathbinrelspacing</code>	76	<code>\Umathopeninnerspacing</code>	76
<code>\Umathchar</code>	66, 87	<code>\Umathopenopspacing</code>	76
<code>\Umathcharclass</code>	85	<code>\Umathopenopspacing</code>	76
<code>\Umathchardef</code>	66, 87	<code>\Umathopenordspacing</code>	76
<code>\Umathcharfam</code>	85	<code>\Umathopenpunctspacing</code>	76
<code>\Umathcharnum</code>	66	<code>\Umathopenrelspacing</code>	76
<code>\Umathcharnumdef</code>	65, 66	<code>\Umathoperatorsize</code>	66, 70, 75
<code>\Umathcharslot</code>	85	<code>\Umathopinnerspacing</code>	76
<code>\Umathclosebinspacing</code>	76	<code>\Umathopopenspacing</code>	76
<code>\Umathcloseclosespacing</code>	76	<code>\Umathopopspacing</code>	76
<code>\Umathcloseinnerspacing</code>	76	<code>\Umathopordspacing</code>	76
<code>\Umathcloseopspacing</code>	76	<code>\Umathoppunctspacing</code>	76
<code>\Umathcloseopspacing</code>	76	<code>\Umathoprelspacing</code>	76
<code>\Umathcloseordspacing</code>	76	<code>\Umathordbinspacing</code>	76
<code>\Umathclosepunctspacing</code>	76	<code>\Umathordclosepacing</code>	76
<code>\Umathcloserelspacing</code>	76	<code>\Umathordinnerspacing</code>	76
<code>\Umathcode</code>	66, 85	<code>\Umathordopspacing</code>	76
<code>\Umathcodenum</code>	66	<code>\Umathordopspacing</code>	76
<code>\Umathconnectoroverlapmin</code>	71, 74	<code>\Umathordordspacing</code>	76
		<code>\Umathordpunctspacing</code>	76



<code>\Umathordrelspacing</code>	76	<code>\Umathunderbarvgap</code>	70
<code>\Umathoverbarkern</code>	70	<code>\Umathunderdelimiterbgap</code>	70, 83
<code>\Umathoverbarrule</code>	70	<code>\Umathunderdelimitervgap</code>	70, 83
<code>\Umathoverbarvgap</code>	70	<code>\Umath*</code>	69
<code>\Umathoverdelimiterbgap</code>	70, 83	<code>\Umath...spacing</code>	76
<code>\Umathoverdelimitervgap</code>	70, 83	<code>\Umiddle</code>	84
<code>\Umathpunctbinspacing</code>	76	<code>\Unosubscript</code>	86
<code>\Umathpunctclosespacing</code>	76	<code>\Unosuperscript</code>	86
<code>\Umathpunctinnerspacing</code>	76	<code>\Uoverdelimiter</code>	66, 83
<code>\Umathpunctopenspacing</code>	76	<code>\Uradical</code>	66, 82
<code>\Umathpuncttopspacing</code>	76	<code>\Uright</code>	84
<code>\Umathpunctordspacing</code>	76	<code>\Uroot</code>	66, 82, 102
<code>\Umathpunctpunctspacing</code>	76	<code>\Uskewed</code>	84
<code>\Umathpunctrelspacing</code>	76	<code>\Uskewedwithdelims</code>	84
<code>\Umathquad</code>	70, 74	<code>\Ustack</code>	68
<code>\Umathradicaldegreeafter</code>	70, 74, 82	<code>\Ustartdisplaymath</code>	86
<code>\Umathradicaldegreebefore</code>	70, 74, 82	<code>\Ustartmath</code>	86
<code>\Umathradicaldegreeraise</code>	70, 74, 75, 82	<code>\Ustopdisplaymath</code>	86
<code>\Umathradicalkern</code>	70	<code>\Ustopmath</code>	86
<code>\Umathradicalrule</code>	70, 74	<code>\Usubscript</code>	86
<code>\Umathradicalvgap</code>	70, 74	<code>\Usuperscript</code>	86
<code>\Umathrelbinspacing</code>	76	<code>\Uunderdelimiter</code>	66, 83
<code>\Umathrelclosespacing</code>	76	<code>\uccode</code>	22, 31, 148
<code>\Umathrelinnerspacing</code>	76	<code>\uchyph</code>	96
<code>\Umathreloppenspacing</code>	76	<code>\unexpanded</code>	168
<code>\Umathreloppspacing</code>	76	<code>\unhbox</code>	31
<code>\Umathrelordspacing</code>	76	<code>\unhcopy</code>	31
<code>\Umathrelpunctspacing</code>	76	<code>\unvbox</code>	31
<code>\Umathrelrelspacing</code>	76	<code>\unvcopy</code>	31
<code>\Umathskewedfractionhgap</code>	84	<code>\uppercase</code>	39
<code>\Umathskewedfractionvgap</code>	84	<code>\useboxresource</code>	47
<code>\Umathspaceafterscript</code>	71, 74	<code>\useimageresource</code>	47
<code>\Umathstackdenomdown</code>	70		
<code>\Umathstacknumup</code>	70	<code>\vadjust</code>	93, 130, 152
<code>\Umathstackvgap</code>	70	<code>\valign</code>	131
<code>\Umathsubshiftdown</code>	70, 79	<code>\vbox</code>	12, 32, 46, 132, 149, 164
<code>\Umathsubshiftdrop</code>	70	<code>\vcenter</code>	132
<code>\Umathsubsupshiftdown</code>	70, 80	<code>\vpack</code>	46
<code>\Umathsubsupvgap</code>	71	<code>\vpack</code>	48
<code>\Umathsubtopmax</code>	71	<code>\vrule</code>	12
<code>\Umathsupbottommin</code>	71	<code>\vskip</code>	12, 94
<code>\Umathsupshiftdrop</code>	70	<code>\vsplit</code>	31, 46, 132, 150
<code>\Umathsupshiftup</code>	70, 80	<code>\vtop</code>	12, 46, 132, 149
<code>\Umathsupsubbottommax</code>	71		
<code>\Umathunderbarkern</code>	70	<code>\wd</code>	31
<code>\Umathunderbarrule</code>	70	<code>\widowpenalties</code>	163



\wordboundary 48, 97

\write 128

\xtoksapp 38

\xtokspre 38

\- 93



Callbacks

b

buildpage_filter 130
build_page_insert 131

c

contribute_filter 130

d

define_font 137

f

find_data_file 128
find_format_file 128
find_read_file 128, 129

h

hpack_filter 131, 133
hyphenate 134

k

kerning 135

l

ligaturing 135
linebreak_filter 132, 163

m

mlist_to_hlist 80, 116, 135

o

open_read_file 128

p

post_linebreak_filter 133
pre_dump 136
pre_linebreak_filter 131, 163
process_jobname 129
process_rule 134

s

show_error_hook 136
show_error_message 137
show_lua_error_hook 137

start_file 137

start_run 136

stop_file 137

stop_run 136

v

vpack_filter 131, 133

w

wrapup_run 137





Nodes

This register contains the nodes that are known to L^AT_EX. The primary nodes are in bold, whatsits that are determined by their subtype are normal. The names prefixed by pdf_ are backend specific.

a

accent 100
adjust 93
attr 111
attribute_list 111

b

boundary 48, 97

c

choice 101

d

delim 100
delta 155
dir 12, 98
disc 12, 32, 93

f

fence 102
fraction 82, 102

g

glue 12, 32, 94
glue-spec 147
glue_spec 94, 95, 143, 145, 147
glyph 12, 32, 96, 109

h

hlist 12, 32, 34, 91, 109

i

ins 92

k

kern 12, 32, 95

l

local_par 97, 163

m

marginkern 98
mark 92
math 94
math_char 99
math_text_char 99

n

noad 100

p

parameter 101
penalty 95

r

radical 101
rule 12, 91

s

style 101
sub_box 99
sub_mlist 99

t

temp 90

v

vlist 12, 32, 91, 109





Libraries

This register contains the functions available in libraries. Not all functions are documented, for instance because they can be experimental or obsolete.

- char_depth 179
- char_height 179
- char_width 179
- fields 175
- get_boolean 179
- get_numeric 179
- get_path 179
- get_string 179
- pen_info 178
- callback**
 - find 127
 - list 127
 - register 127
- lua**
 - bytecode 139
 - getbytecode 139
 - getcurrenttime 140
 - getluaname 140
 - getpreciseseconds 140
 - getpreciseticks 140
 - getruntime 140
 - getstacktop 140
 - name 140
 - newindex 139
 - newtable 139
 - setbytecode 139
 - setluaname 140
 - startupfile 139
 - version 139
- mplib**
 - execute 174
 - finish 175
 - new 173
 - statistics 174
 - version 173
- node**
 - check_discretionaries 114
 - check_discretionary 114
 - copy 106, 118
 - copy_list 106, 118
 - count 107
 - current_attr 111, 118
 - dimensions 115
 - end_of_math 116
 - fields 89, 105
 - find_attribute 112
 - find_node 109
 - first_glyph 112
 - flatten_discretionaries 114
 - flush_list 105, 119
 - flush_node 105, 119
 - flush_properties_table 122
 - free 105, 119
 - getfield 119
 - getglue 110
 - getproperty 119
 - get_attribute 112, 119
 - get_properties_table 119, 122
 - has_attribute 112, 119
 - has_field 105, 119
 - has_glyph 113
 - hpack 114
 - id 104
 - insert_after 108, 120
 - insert_before 107, 120
 - is_char 113
 - is_glyph 113
 - is_node 105, 120
 - is_zero_glue 110
 - kerning 113
 - last_node 108
 - length 107
 - ligaturing 113
 - mlist_to_hlist 116
 - new 105, 120
 - prepend_prevdepth 115
 - protect_glyph 113
 - protect_glyphs 113
 - protrusion_skippable 114
 - rangedimensions 115



- remove 107, 120
- setfield 120
- setglue 110, 120
- setproperty 121
- set_attribute 112, 120
- set_properties_mode 122
- slide 106
- subtypes 89
- tail 107, 121
- todirect 117
- tonode 117
- tostring 117, 121
- traverse 108, 121
- traverse_char 109, 121
- traverse_glyph 109, 121
- traverse_id 109, 121
- traverse_list 109, 121
- type 104, 121
- types 104
- unprotect_glyph 113
- unprotect_glyphs 113
- unset_attribute 112, 121
- values 89
- vpack 115
- write 106, 121

node.direct

- check_discretionaries 118
- check_discretionary 118
- copy 118
- copy_list 118
- count 118
- current_attr 118
- dimensions 118
- effective_glue 118
- end_of_math 118
- find_attribute 118
- first_glyph 118
- flatten_discretionaries 118
- flush_list 119
- flush_node 119
- free 119
- getattributelist 119
- getboth 119
- getbox 119
- getchar 119
- getcomponents 119

- getdata 119
- getdepth 119
- getdirection 119
- getdisc 119
- getexpansion 119
- getfam 119
- getfield 119
- getfont 119
- getglue 119
- getheight 119
- getid 119
- getkern 119
- getlang 119
- getleader 119
- getlist 119
- getnext 119
- getnormalizedline 119
- getnucleus 119
- getoffsets 119
- getorientation 119
- getpenalty 119
- getpost 119
- getpre 119
- getprev 119
- getproperty 119
- getreplace 119
- getshift 119
- getsub 119
- getsubtype 119
- getsup 119
- getwhd 119
- getwidth 119
- get_attribute 119
- get_properties_table 119
- get_synctex_fields 119
- has_attribute 119
- has_dimensions 119
- has_field 119
- has_glyph 120
- hpack 120
- insert_after 120
- insert_before 120
- is_char 120
- is_direct 120
- is_glyph 120
- is_node 120



- is_valid 120
- is_zero_glue 120
- kerning 120
- last_node 120
- length 120
- ligaturing 120
- make_extensible 120
- mlist_to_hlist 120
- naturalwidth 120
- new 120
- prepend_prevdepth 120
- protect_glyph 120
- protect_glyphs 120
- protrusion_skippable 120
- rangedimensions 120
- remove 120
- setattributelist 120
- setboth 120
- setbox 120
- setchar 120
- setcomponents 120
- setdata 120
- setdepth 120
- setdirection 120
- setdisc 120
- setexpansion 120
- setfam 120
- setfield 120
- setfont 120
- setglue 120
- setheight 120
- setkern 120
- setlang 120
- setleader 120
- setlink 120
- setlist 120
- setnext 121
- setnucleus 121
- setoffsets 121
- setorientation 121
- setpenalty 121
- setprev 121
- setproperty 121
- setshift 121
- setsplit 121
- setsub 121

- setsubtype 121
- setup 121
- setwhd 121
- setwidth 121
- set_attribute 120
- set_synctex_fields 120
- slide 121
- start_of_par 121
- tail 121
- todirect 121
- tonode 121
- traverse 121
- traverse_char 121
- traverse_glyph 121
- traverse_id 121
- traverse_list 121
- unprotect_glyph 121
- unprotect_glyphs 121
- unset_attribute 121
- usedlist 121
- uses_font 121
- vpack 121
- write 121

os

- env 193
- gettimeofday 193
- name 193
- selfarg 193
- selfdir 193
- setenv 193
- type 193
- uname 193

pdfc

- arraytotable 184
- close 181
- closestream 183
- dictionarytotable 184
- getarray 183
- getboolean 183
- getbox 182
- getcatalag 182
- getdictionary 183
- getfromarray 183, 184
- getfromdictionary 183, 184
- getfromreference 184
- getfromstream 183



getinfo 182
getinteger 183
getname 183
getnofobjects 182
getnofpages 182
getnumber 183
getpage 182
getsize 182
getstatus 181
getstream 183
getstring 183
gettrailer 182
getversion 182
new 181, 185
open 181
openstream 183
readfromstream 183
readfromwholestream 183
unencrypt 181

sio

getposition 187
readbytes 187
readbytetable 187
readcardinaltable 187
readcardinal1 187
readcardinal2 187
readcardinal3 187
readcardinal4 187
readfixed2 187
readfixed4 187
readintegertable 187
readinteger1 187
readinteger2 187
readinteger3 187
readinteger4 187
read2dot14 187
setposition 187
skipposition 187

status

list 140
resetmessages 140
setexitcode 140

string

bytepairs 192
bytes 192
characterpairs 192

characters 192
explode 192
utfcharacter 193
utfcharacters 192
utflength 193
utfvalue 193
utfvalues 192

tex

attribute 146
badness 162
box 146, 149
catcode 148
count 146
cprint 154
definefont 158
delcode 148
dimen 146
enableprimitives 158
error 156
extraprimitives 159
fontidentifier 156
fontname 156
forcehmode 158
force_synctex_line 164
force_synctex_tag 164
get 142
getattribute 146
getbox 146, 149
getcatcode 148
getcount 146
getdelcode 148
getdelcodes 148
getdimen 146
getfamilyoffont 157
getglue 146
gethelptext 156
getinteraction 157
getlccode 148
getlinenumber 156
getlist 151
getlocallevel 164
getmath 150
getmathcode 148
getmathcodes 148
getmuglue 146
getmuskip 146



- getnest 152
- getpagestate 164
- getsfcode 148
- getskip 146
- gettoks 146
- getuccode 148
- get_synctex_line 164
- get_synctex_mode 164
- get_synctex_tag 164
- glue 146
- hashtokens 158
- init_rand 164
- isattribute 146
- isbox 146
- iscount 146
- isdimen 146
- isglue 146
- ismuglue 146
- ismuskip 146
- isskip 146
- istoks 146
- lccode 148
- linebreak 163
- lists 151
- lua_math_random 164
- lua_math_randomseed 164
- mathcode 148
- muglue 146
- muskip 146
- nest 152
- normal_rand 164
- number 155
- primitives 162
- print 153
- ptr 152
- resetparagraph 162
- romannumeral 155
- round 155
- scale 155
- scantoks 146
- set 142
- setattribute 146
- setbox 146, 149
- setcatcode 148
- setcount 146
- setdelcode 148

- setdelcodes 148
- setdimen 146
- setglue 146
- setinteraction 157
- setlccode 148
- setlinenumber 156
- setlist 151
- setmath 150
- setmathcode 148
- setmathcodes 148
- setmuglue 146
- setmuskip 146
- setsfcode 148
- setskip 146
- settoks 146
- setuccode 148
- set_synctex_line 164
- set_synctex_mode 164
- set_synctex_no_files 164
- set_synctex_tag 164
- sfcode 148
- shipout 164
- show_context 156
- skip 146
- sp 156
- splitbox 150
- sprint 153
- toks 146
- tprint 154
- triggerbuildpage 150
- uccode 148
- uniformdeviate 164
- uniform_rand 164
- write 155

texio

- closeinput 166
- setescape 166
- write 166
- write_nl 166

token

- biggest_char 169
- commands 169
- command_id 169
- create 169
- expand 169
- get_active 169



get_cmdname 169
get_command 169
get_csname 169
get_expandable 169
get_functions_table 170
get_id 169
get_index 169
get_macro 170
get_meaning 170
get_mode 169
get_next 169, 171
get_protected 169
get_tok 169
is_defined 169
is_token 169
new 169
put_next 171
scan_argument 166
scan_code 166
scan_csname 166
scan_dimen 166
scan_float 166
scan_glue 166
scan_int 166
scan_keyword 166
scan_keywordcs 166
scan_list 166
scan_real 166
scan_string 166
scan_token 169
scan_toks 166
scan_word 166
set_char 170
set_lua 170
set_macro 170



Statistics

The following fonts are used in this document:

used	filesize	version	filename
22	988.684	5.000	cambmth.ttf
4	927.280	5.020	cambria.ttf
11	163.452	1.802	LucidaBrightMathOT-Demi.otf
11	348.296	1.802	LucidaBrightMathOT.otf
4	73.284	1.801	LucidaBrightOT.otf
22	733.500	1.958	latinmodern-math.otf
1	64.684	2.004	lmmono10-regular.otf
1	64.160	2.004	lmmonoltcond10-regular.otf
4	111.536	2.004	lmroman10-regular.otf
22	525.008	1.106	texgyredejavu-math.otf
22	601.220	1.632	texgyrepagella-math.otf
4	218.100	2.501	texgyrepagella-regular.otf
1	693.876	2.340	DejaVuSans-Bold.ttf
1	741.536	2.340	DejaVuSans.ttf
5	318.392	2.340	DejaVuSansMono-Bold.ttf
1	245.948	2.340	DejaVuSansMono-Oblique.ttf
4	335.068	2.340	DejaVuSansMono.ttf
13	345.364	2.340	DejaVuSerif-Bold.ttf
1	336.884	2.340	DejaVuSerif-BoldItalic.ttf
2	343.388	2.340	DejaVuSerif-Italic.ttf
5	367.260	2.340	DejaVuSerif.ttf
161	8.546.920		21 files loaded



