


LuaTeX

Reference



March 2016
Version 0.90.0

LuaTeX

Reference

Manual

copyright : LuaTeX development team
more info : www.luatex.org
version : March 25, 2016

Contents

Introduction	15
1 Basic T_EX enhancements	17
1.1 Introduction	17
1.2 Version information	17
1.3 UNICODE text support	18
1.4 Extended tables	18
1.5 Attributes	19
1.5.1 Attribute registers	19
1.5.2 Box attributes	19
1.6 LUA related primitives	20
1.6.1 <code>\directlua</code>	20
1.6.2 <code>\latelua</code>	21
1.6.3 <code>\luaescapestring</code>	22
1.6.4 <code>\luafunction</code>	22
1.7 <code>\clearmarks</code>	22
1.8 <code>\noligs</code> and <code>\nokerns</code>	23
1.9 <code>\formatname</code>	23
1.10 <code>\scantextokens</code>	23
1.11 Alignments	23
1.11.1 <code>\alignmark</code>	23
1.11.2 <code>\aligntab</code>	23
1.12 Catcode tables	23
1.12.1 <code>\catcodetable</code>	24
1.12.2 <code>\initcatcodetable</code>	24
1.12.3 <code>\savecatcodetable</code>	24
1.13 Suppressing errors	24
1.13.1 <code>\suppressfontnotfounderror</code>	24
1.13.2 <code>\suppresslongerror</code>	25
1.13.3 <code>\suppressifcsnameerror</code>	25
1.13.4 <code>\suppressoutererror</code>	25
1.13.5 <code>\suppressmathparerror</code>	25



1.14	<code>\matheqnogapstep</code>	25
1.15	<code>\outputbox</code>	26
1.16	<code>\fontid</code> and <code>\setfontid</code>	26
1.17	<code>\gleaders</code>	26
1.18	<code>\nohrule</code> and <code>\novrule</code>	26
1.19	<code>\Uchar</code>	26
1.20	<code>\hyphenationmin</code>	27
1.21	<code>\boundary</code> and <code>\noboundary</code>	27
1.22	Debugging	27
1.23	Images and Forms	27
1.24	<code>\outputmode</code> and <code>\draftmode</code>	28
1.25	File syntax	28
1.26	Font syntax	28
1.27	Writing to file	28
1.28	<code>\nospaces</code>	28
1.29	<code>\letcharcode</code>	29
2	LUA general	31
2.1	Initialization	31
2.1.1	LUAT _E X as a Lua interpreter	31
2.1.2	LUAT _E X as a Lua byte compiler	31
2.1.3	Other commandline processing	31
2.2	Lua behaviour	34
2.3	Lua modules	37
3	Languages and characters, fonts and glyphs	39
3.1	Characters and glyphs	39
3.2	The main control loop	41
3.3	Loading patterns and exceptions	42
3.4	Applying hyphenation	43
3.5	Applying ligatures and kerning	44
3.6	Breaking paragraphs into lines	46
4	Font structure	49
4.1	Real fonts	54
4.2	Virtual fonts	56
4.2.1	Artificial fonts	58
4.2.2	Example virtual font	58



5	Math	59
5.1	The current math style	59
5.1.1	<code>\mathstyle</code>	59
5.1.2	<code>\Ustack</code>	59
5.2	Unicode math characters	59
5.3	Cramped math styles	61
5.4	Math parameter settings	62
5.5	Skips around display math	64
5.6	Font-based Math Parameters	64
5.7	Math spacing setting	66
5.8	Math accent handling	67
5.9	Math root extension	68
5.10	Math kerning in super- and subscripts	68
5.11	Scripts on horizontally extensible items like arrows	69
5.12	Extracting values	70
5.13	fractions	70
5.14	Other Math changes	71
5.14.1	Verbose versions of single-character math commands	71
5.14.2	Allowed math commands in non-math modes	71
5.15	Math surrounding skips	71
5.15.1	Delimiters: <code>\Uleft</code> , <code>\Umiddle</code> and <code>\Uright</code>	72
5.15.2	Fixed scripts	72
5.15.3	Tracing	73
5.15.4	Math options	73
5.15.4.1	<code>\mathoption noitaliccompensation</code>	73
5.15.4.2	<code>\mathoption nocharitalic</code>	73
5.15.4.3	<code>\mathoption useoldfractionscaling</code>	74
6	Nodes	75
6.1	LUA node representation	75
6.1.1	Auxiliary items	75
6.1.1.1	<code>attribute_list</code> and attribute items	75
6.1.2	Main text nodes	76
6.1.2.1	<code>hlist</code> nodes	76
6.1.2.2	<code>vlist</code> nodes	77



6.1.2.3	rule nodes	77
6.1.2.3.1	normal rules	77
6.1.2.4	ins nodes	77
6.1.2.5	mark nodes	78
6.1.2.6	adjust nodes	78
6.1.2.7	disc nodes	78
6.1.2.8	math nodes	79
6.1.2.9	glue nodes	79
6.1.2.10	kern nodes	81
6.1.2.11	penalty nodes	81
6.1.2.12	glyph nodes	81
6.1.2.13	margin_kern nodes	82
6.1.3	Math nodes	82
6.1.3.1	Math kernel subnodes	83
6.1.3.1.1	math_char and math_text_char subnodes	83
6.1.3.1.2	sub_box and sub_mlist subnodes	83
6.1.3.2	Math delimiter subnode	83
6.1.3.2.1	delim subnodes	83
6.1.3.3	Math core nodes	84
6.1.3.3.1	simple nodes	84
6.1.3.3.2	accent nodes	85
6.1.3.3.3	style nodes	85
6.1.3.3.4	choice nodes	85
6.1.3.3.5	radical nodes	85
6.1.3.3.6	fraction nodes	86
6.1.3.3.7	fence nodes	86
6.1.4	whatsit nodes	87
6.1.4.1	open nodes	87
6.1.4.2	write nodes	87
6.1.4.3	close nodes	87
6.1.4.4	special nodes	87
6.1.4.5	boundary nodes	88
6.1.4.6	language nodes	88
6.1.4.7	local_par nodes	88



6.1.4.8	dir nodes	88
6.1.4.9	pdf_literal nodes	89
6.1.4.10	pdf_refobj nodes	89
6.1.4.11	pdf_annot nodes	89
6.1.4.12	pdf_start_link nodes	90
6.1.4.13	pdf_end_link nodes	90
6.1.4.14	pdf_dest nodes	90
6.1.4.15	pdf_action nodes	90
6.1.4.16	pdf_thread nodes	91
6.1.4.17	pdf_start_thread nodes	91
6.1.4.18	pdf_end_thread nodes	91
6.1.4.19	save_pos nodes	91
6.1.4.20	late_lua nodes	92
6.1.4.21	pdf_colorstack nodes	92
6.1.4.22	pdf_setmatrix nodes	92
6.1.4.23	pdf_save nodes	92
6.1.4.24	pdf_restore nodes	92
6.1.4.25	user_defined nodes	93
6.2	Two access models	93
7	LUAT_EX LUA Libraries	97
7.1	The callback library	97
7.1.1	File discovery callbacks	98
7.1.1.1	find_read_file and find_write_file	98
7.1.1.2	find_font_file	98
7.1.1.3	find_output_file	98
7.1.1.4	find_format_file	99
7.1.1.5	find_vf_file	99
7.1.1.6	find_map_file	99
7.1.1.7	find_enc_file	99
7.1.1.8	find_sfd_file	99
7.1.1.9	find_pk_file	99
7.1.1.10	find_data_file	99
7.1.1.11	find_opentype_file	99
7.1.1.12	find_truetype_file and find_type1_file	100



7.1.1.13	find_image_file	100
7.1.2	File reading callbacks	100
7.1.2.1	open_read_file	100
7.1.2.1.1	reader	101
7.1.2.1.2	close	101
7.1.2.2	General file readers	101
7.1.3	Data processing callbacks	102
7.1.3.1	process_input_buffer	102
7.1.3.2	process_output_buffer	102
7.1.3.3	process_jobname	102
7.1.4	Node list processing callbacks	103
7.1.4.1	contribute_filter	103
7.1.4.2	buildpage_filter	103
7.1.4.3	pre_linebreak_filter	103
7.1.4.4	linebreak_filter	104
7.1.4.5	append_to_vlist_filter	105
7.1.4.6	post_linebreak_filter	105
7.1.4.7	hpack_filter	105
7.1.4.8	vpack_filter	105
7.1.4.9	hpack_quality	106
7.1.4.10	vpack_quality	106
7.1.4.11	process_rule	106
7.1.4.12	pre_output_filter	106
7.1.4.13	hyphenate	107
7.1.4.14	ligaturing	107
7.1.4.15	kerning	107
7.1.4.16	insert_local_par	108
7.1.4.17	mlist_to_hlist	108
7.1.5	Information reporting callbacks	108
7.1.5.1	pre_dump	108
7.1.5.2	start_run	108
7.1.5.3	stop_run	108
7.1.5.4	start_page_number	109
7.1.5.5	stop_page_number	109



7.1.5.6	show_error_hook	109
7.1.5.7	show_error_message	109
7.1.5.8	show_lua_error_hook	109
7.1.5.9	start_file	109
7.1.5.10	stop_file	110
7.1.6	PDF-related callbacks	110
7.1.6.1	finish_pdffile	110
7.1.6.2	finish_pdfpage	110
7.1.7	Font-related callbacks	110
7.1.7.1	define_font	110
7.2	The epdf library	111
7.3	The font library	119
7.3.1	Loading a TFM file	120
7.3.2	Loading a VF file	120
7.3.3	The fonts array	120
7.3.4	Checking a font's status	120
7.3.5	Defining a font directly	121
7.3.6	Projected next font id	121
7.3.7	Font id	121
7.3.8	Currently active font	121
7.3.9	Maximum font id	121
7.3.10	Iterating over all fonts	121
7.4	The fontloader library	122
7.4.1	Getting quick information on a font	122
7.4.2	Loading an OPENTYPE or TRUETYPE file	122
7.4.3	Applying a 'feature file'	123
7.4.4	Applying an 'AFM file'	124
7.4.5	Fontloader font tables	124
7.4.5.1	Table types	125
7.4.5.1.1	Top-level	125
7.4.5.1.2	Glyph items	127
7.4.5.1.3	map table	129
7.4.5.1.4	private table	130
7.4.5.1.5	cidinfo table	130



7.4.5.1.6	pfminfo table	131
7.4.5.1.7	names table	132
7.4.5.1.8	anchor_classes table	133
7.4.5.1.9	gpos table	133
7.4.5.1.10	gsub table	134
7.4.5.1.11	ttf_tables and ttf_tab_saved tables	134
7.4.5.1.12	mm table	134
7.4.5.1.13	mark_classes table	135
7.4.5.1.14	math table	135
7.4.5.1.15	validation_state table	136
7.4.5.1.16	horiz_base and vert_base table	137
7.4.5.1.17	altuni table	137
7.4.5.1.18	vert_variants and horiz_variants table	137
7.4.5.1.19	mathkern table	138
7.4.5.1.20	kerns table	138
7.4.5.1.21	vkerns table	138
7.4.5.1.22	texdata table	138
7.4.5.1.23	lookups table	138
7.5	The img library	139
7.5.1	img.new	139
7.5.2	img.keys	140
7.5.3	img.scan	142
7.5.4	img.copy	142
7.5.5	img.write	142
7.5.6	img.immediatewrite	143
7.5.7	img.node	143
7.5.8	img.types	143
7.5.9	img.bboxes	143
7.6	The kpse library	144
7.6.1	kpse.set_program_name and kpse.new	144
7.6.2	find_file	144
7.6.3	lookup	145
7.6.4	init_prog	146
7.6.5	readable_file	146



7.6.6	expand_path	146
7.6.7	expand_var	146
7.6.8	expand_braces	147
7.6.9	show_path	147
7.6.10	var_value	147
7.6.11	version	147
7.7	The lang library	147
7.8	The lua library	149
7.8.1	LUA bytecode registers	149
7.8.2	LUA chunk name registers	150
7.9	The mplib library	150
7.9.1	mplib.new	150
7.9.2	mp:statistics	151
7.9.3	mp:execute	151
7.9.4	mp:finish	151
7.9.5	Result table	152
7.9.5.1	fill	153
7.9.5.2	outline	153
7.9.5.3	text	153
7.9.5.4	special	154
7.9.5.5	start_bounds, start_clip	154
7.9.5.6	stop_bounds, stop_clip	154
7.9.6	Subsidiary table formats	154
7.9.6.1	Paths and pens	154
7.9.6.2	Colors	154
7.9.6.3	Transforms	155
7.9.6.4	Dashes	155
7.9.7	Character size information	155
7.9.7.1	mp:char_width	155
7.9.7.2	mp:char_height	155
7.9.7.3	mp:char_depth	155
7.10	The node library	155
7.10.1	Node handling functions	156
7.10.1.1	node.is_node	156



7.10.1.2	<code>node.types</code>	156
7.10.1.3	<code>node.whatsits</code>	156
7.10.1.4	<code>node.id</code>	157
7.10.1.5	<code>node.subtype</code>	157
7.10.1.6	<code>node.type</code>	157
7.10.1.7	<code>node.fields</code>	157
7.10.1.8	<code>node.has_field</code>	157
7.10.1.9	<code>node.new</code>	157
7.10.1.10	<code>node.free</code>	158
7.10.1.11	<code>node.flush_list</code>	158
7.10.1.12	<code>node.copy</code>	158
7.10.1.13	<code>node.copy_list</code>	158
7.10.1.14	<code>node.next</code>	158
7.10.1.15	<code>node.prev</code>	159
7.10.1.16	<code>node.current_attr</code>	159
7.10.1.17	<code>node.hpack</code>	159
7.10.1.18	<code>node.vpack</code>	160
7.10.1.19	<code>node.dimensions</code>	160
7.10.1.20	<code>node.mlist_to_hlist</code>	161
7.10.1.21	<code>node.slide</code>	161
7.10.1.22	<code>node.tail</code>	161
7.10.1.23	<code>node.length</code>	161
7.10.1.24	<code>node.count</code>	161
7.10.1.25	<code>node.traverse</code>	162
7.10.1.26	<code>node.traverse_id</code>	162
7.10.1.27	<code>node.end_of_math</code>	163
7.10.1.28	<code>node.remove</code>	163
7.10.1.29	<code>node.insert_before</code>	163
7.10.1.30	<code>node.insert_after</code>	163
7.10.1.31	<code>node.first_glyph</code>	164
7.10.1.32	<code>node.ligaturing</code>	164
7.10.1.33	<code>node.kerning</code>	164
7.10.1.34	<code>node.unprotect_glyphs</code>	164
7.10.1.35	<code>node.protect_glyphs</code>	164



7.10.1.36	node.last_node	164
7.10.1.37	node.write	165
7.10.1.38	node.protrusion_skipable	165
7.10.2	Attribute handling	165
7.10.2.1	node.has_attribute	165
7.10.2.2	node.set_attribute	165
7.10.2.3	node.unset_attribute	165
7.11	The pdf library	166
7.11.1	pdf.mapfile, pdf.mapline	166
7.11.2	pdf.catalog, pdf.info, pdf.names, pdf.trailer	166
7.11.3	pdf.<set/get>pageattributes, pdf.<set/get>pageresources, pdf.<set/get>pagesattributes	166
7.11.4	pdf.<set/get>xformattributes, pdf.<set/get>xformresources	166
7.11.5	pdf.setcompresslevel and pdf.setobjcompresslevel	167
7.11.6	pdf.setdecimaldigits and pdf.getdecimaldigits	167
7.11.7	pdf.setpkresolution and pdf.getpkresolution	167
7.11.8	pdf.lastobj, pdf.lastlink, pdf.lastannot, and pdf.retval	167
7.11.9	pdf.setorigin, pdf.getorigin	167
7.11.10	pdf.setlinkmargin, pdf.getlinkmargin pdf.setdestmargin, pdf.getdestmargin pdf.setthreadmargin, pdf.getthreadmargin pdf.setxformmargin, pdf.getxformmargin	167
7.11.11	pdf.h, pdf.v	167
7.11.12	pdf.getpos, pdf.getpos, pdf.getvpos	167
7.11.13	pdf.hasmatrix, pdf.getmatrix	168
7.11.14	pdf.print	168
7.11.15	pdf.immediateobj	168
7.11.16	pdf.obj	169
7.11.17	pdf.refobj	170
7.11.18	pdf.reserveobj	170
7.11.19	pdf.registerannot	170
7.11.20	pdf.newcolorstack	170
7.12	The pdfscanner library	170
7.13	The status library	173
7.14	The tex library	175
7.14.1	Internal parameter values	175



7.14.1.1	Integer parameters	176
7.14.1.2	Dimension parameters	177
7.14.1.3	Direction parameters	178
7.14.1.4	Glue parameters	178
7.14.1.5	Muglue parameters	178
7.14.1.6	Tokenlist parameters	178
7.14.2	Convert commands	178
7.14.3	Last item commands	179
7.14.4	Attribute, count, dimension, skip and token registers	179
7.14.5	Character code registers	180
7.14.6	Box registers	181
7.14.7	Math parameters	182
7.14.8	Special list heads	184
7.14.9	Semantic nest levels	184
7.14.10	Print functions	185
7.14.10.1	tex.print	185
7.14.10.2	tex.sprint	185
7.14.10.3	tex.tprint	186
7.14.10.4	tex.cprint	186
7.14.10.5	tex.write	186
7.14.11	Helper functions	187
7.14.11.1	tex.round	187
7.14.11.2	tex.scale	187
7.14.11.3	tex.sp	187
7.14.11.4	tex.definefont	187
7.14.11.5	tex.getlinenumber and tex.setlinenumber	188
7.14.11.6	tex.error	188
7.14.11.7	tex.hashtokens	188
7.14.12	Functions for dealing with primitives	188
7.14.12.1	tex.enableprimitives	188
7.14.12.2	tex.extraprimitives	189
7.14.12.3	tex.primitives	192
7.14.13	Core functionality interfaces	192
7.14.13.1	tex.badness	192



7.14.13.2 <code>tex.resetparagraph</code>	192
7.14.13.3 <code>tex.linebreak</code>	192
7.14.13.4 <code>tex.shipout</code>	194
7.15 The <code>texconfig</code> table	194
7.16 The <code>texio</code> library	195
7.16.1 Printing functions	195
7.16.1.1 <code>texio.write</code>	195
7.16.1.2 <code>texio.write_nl</code>	195
7.16.1.3 <code>texio.setescape</code>	196
7.16.2 The token libray	196
8 Modifications	201
8.1 The merged engines	201
8.1.1 The need for change	201
8.1.2 Changes from T _E X 3.1415926	201
8.1.3 Changes from ϵ -T _E X 2.2	201
8.1.4 Changes from PDFT _E X 1.40	202
8.1.5 Changes from ALEPH RC4	204
8.1.6 Changes from standard WEB2C	205
8.1.7 The backend primitives <code>\pdf *</code>	205
8.2 Implementation notes	210
8.2.1 Memory allocation	210
8.2.2 Sparse arrays	211
8.2.3 Simple single-character csnames	211
8.2.4 Compressed format	211
8.2.5 Binary file reading	211





Introduction

This book will eventually become the reference manual of Lua \TeX . At the moment, it simply reports the behaviour of the executable matching the snapshot or beta release date in the title page. We don't claim it is complete and we assume that the reader knows about \TeX as described in "The \TeX Book", the " ϵ - \TeX manual", the "pdf \TeX manual", etc. Additional reference material is published in journals of user groups and Con \TeX t related documentation.

Successive versions brought new functionality, more control, some cleanup of internals and experimental features evolved into stable ones or were dropped. Already quite early Lua \TeX could be used for production and it was used in production by the authors. Successive versions sometimes demanded a adaption to the Lua interfacing, but the concepts were unchanged. The current version can be considered stable in functionality and there will be no fundamental changes between 0.90 and 1.00. Of course we then can decide to move towards version 2.00 with different properties.

Don't expect Lua \TeX to behave the same as pdf \TeX ! Although the core functionality of that 8 bit engine was starting point, it has been combined with the directional support of Omega (Aleph). But, Lua \TeX can behave different due to its wide (32 bit) characters, many registers and large memory support. There is native utf input, support for large (more that 8 bit) fonts, and the math machinery is tuned for OpenType math. There is support for directional typesetting too. The log output can differ from other engines and will likely differ more as we move forward. When you run plain \TeX for sure Lua \TeX runs slower than pdf \TeX but when you run for instance Con \TeX t MkIV it might be faster on more complex documents. But in any case: 32 bit all-over combined with more features has a price.

Lua \TeX consists of a number of interrelated but (still) distinguishable parts. The organization of the source code is adapted so that it can glue all these components together. We continue cleaning up side effects of the accumulated code in \TeX engines (especially code that is not needed any longer).

- Most of pdf \TeX version 1.40.9, converted to C (with patches from later releases). Some experimental features have been removed and some utility macros are not inherited as their functionality can be done in Lua. The number of backend interface commands has been reduced to a few. The extensions are separated from the core (which we keep close to the original \TeX core). Some mechanisms like expansion and protrusion can behave different from the original due to some cleanup and optimization. Some whatsit based functionality (image support and reusable content) is now core functionality.
- The direction model and some other bits from Aleph RC4 (derived from Omega) is included. The related primitives are part of core Lua \TeX but at the node level directional support is no longer based on so called whatsits but on real nodes. In fact, whatsits are now only used for backend specific extensions.
- Neither Aleph's I/O translation processes, nor tcx files, nor enc \TeX can be used, these encoding-related functions are superseded by a Lua-based solution (reader callbacks). In a similar fashion all file io can be intercepted.
- We currently use Lua 5.2.*. At some point we might decide to move to 5.3.* but that is yet to be decided. There are few Lua libraries that we consider part of the core Lua machinery, for instance [lpeg](#). There are additional Lua libraries that interface to the internals of \TeX .



- There are various \TeX extensions but only those that cannot be done using the Lua interfaces. The math machinery often has two code paths: one traditional and the other more suitable for wide OpenType fonts.
- The fontloader uses parts of FontForge 2008.11.17 combined with additional code specific for usage in a \TeX engine. We try to minimize specific font support to what \TeX needs: character references and dimensions and delegate everything else to Lua. That way we keep \TeX open for extensions without touching the core.
- The MetaPost library is integral part of Lua \TeX . This gives \TeX some graphical capabilities using a relative high speed graphical subsystem. Again Lua is used as glue between the frontend and backend. Further development of MetaPost is closely related to Lua \TeX .

The yearly \TeX Live version is the stable version, any version between them are to be considered beta. The beta releases are normally available via the Con \TeX t distribution channels (the garden and so called minimal).

Hans Hagen, Harmut Henkel,
Taco Hoekwater & Luigi Scarso

Version : March 25, 2016
Lua \TeX : Snapshot 90.0
Con \TeX t : 2016.03.17 13:41



1 Basic T_EX enhancements

1.1 Introduction

From day one, LuaT_EX has offered extra features compared to the superset of pdfT_EX and Aleph. That has not been limited to the possibility to execute Lua code via `\directlua`, but LuaT_EX also adds functionality via new T_EX-side primitives.

When LuaT_EX starts up in ‘iniluatex’ mode (`luatex -ini`), it defines only the primitive commands known by T_EX82 and the one extra command `\directlua`. As is fitting, a Lua function has to be called to add the extra primitives to the user environment. The simplest method to get access to all of the new primitive commands is by adding this line to the format generation file:

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

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, you can look up the details in section 7.14.12. 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 pdfT_EX primitives to core LuaT_EX ones, and the ones inherited from Aleph (Omega) are also promoted. Effectively this means that we now have the `tex`, `etex`, `luatex` and `pdftex` (sub)sets left.

1.2 Version information

There are three new primitives to test the version of LuaT_EX:

primitive	explanation	value
<code>\luatexbanner</code>	the banner reported on the command line	This is LuajitTeX, Version beta-0.90.0 (TeX Live 2016)
<code>\luatexversion</code>	a combination of major and minor number	90
<code>\luatexrevision</code>	the revision number, the current value is	0

The official LuaT_EX 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 the given 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.

1.3 UNICODE text support

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 LuaTeX there is no clear separation between the two concepts. Because the subtype of a glyph node can be changed in Lua it is also up to the user.

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 equals sign) or secondary (right of the equals sign) value are: `\char`, `\lccode`, `\uccode`, `\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 a later chapter.

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 $c \geq 1,114,112$, LuaTeX will actually print the single byte corresponding to c minus 1,114,112.

Output to the terminal uses `^^` notation for the lower control range ($c < 32$), with the exception of `^^I`, `^^J` and `^^M`. These are considered ‘safe’ and therefore printed as-is.

Normalization of the Unicode input can be handled by a macro package during callback processing (this will be explained in section 7.1.2).

1.4 Extended tables

All traditional TeX and ϵ -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>	

The glyph properties `\efcode`, `\lpcode` and `\rpcode`, introduced in pdfTeX that deal with font expansion (hz) and character protruding, are also 16-bit. Because font memory management



has been rewritten, these character properties are no longer shared among fonts instances that originate from the same metric file.

1.5 Attributes

1.5.1 Attribute registers

Attributes are a completely new concept in Lua \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 <cname> <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: `-"7FFFFFFF` in hexadecimal, a.k.a. `-2147483647` in decimal. It follows that the value `-"7FFFFFFF` cannot be used as a legal attribute value, but you *can* assign `-"7FFFFFFF` to 'unset' an attribute. All attributes start out in this 'unset' state in `ini \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 6.

1.5.2 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 Lua \TeX 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`. An example:



```

\attribute2=5
\setbox0=\hbox {Hello}
\setbox2=\hbox attr1=12 attr2=-"7FFFFFFF{Hello}

```

This will set the attribute list of box 2 to 1 = 12, and the attributes of box 0 will be 2 = 5. As you can see, assigning the maximum negative value causes an attribute to be ignored.

The `attr` keyword(s) should come before a `to` or `spread`, if that is also specified.

1.6 LUA related primitives

1.6.1 `\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>
\directlua name <general text> <general text>
\directlua <16-bit number> <general text>

```

The last `<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.

The `name <general text>` specifies the name of the Lua chunk, mainly shown in the stack backtrace of error messages created by Lua code. The `<general text>` is expanded fully, thus macros can be used to generate the chunk name, i.e.

```

\directlua name{\jobname:\the\inputlineno} ...

```

to include the name of the input file as well as the input line into the chunk name.

Likewise, the `<16-bit number>` designates a name of a Lua chunk, but in this case the name will be taken from the `lua.name` array (see the documentation of the `lua` table further in this manual).

The chunk name should not start with a `@`, or it will be displayed as a file name (this is a quirk in the current Lua implementation).



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 7.14.10.

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 Lua \TeX 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.

The behaviour documented in the above subsection is considered stable in the sense that there will not be backward-incompatible changes any more.

1.6.2 `\latelua`

`\latelua` stores Lua code in a whatsit that will be processed at the time of shipping out. Its intended use is a cross between `\pdfliteral` and `\write`. Within the Lua code you can print pdf statements directly to the pdf file via `pdf.print`, or you can write to other output streams via `texio.write` or simply using Lua I/O routines.

```
\latelua <general text>
\latelua name <general text> <general text>
\latelua <16-bit number> <general text>
```

Expansion of macros etcetera in the final `<general text>` is delayed until just before the whatsit is executed (like in `\write`). With regard to pdf output stream `\latelua` behaves as `\pdfliteral page`. The name `<general text>` and `<16-bit number>` behave in the same way as they do for `\directlua`



1.6.3 `\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 the \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') }
```

1.6.4 `\luafunction`

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 use this primitive hundreds or thousands of times, it can become noticeable. 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
\uafunction2
```

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
}
```

1.7 `\clearmarks`

This primitive complements the ϵ - \TeX 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>
```



1.8 `\noligs` and `\nokerns`

These primitives prohibit ligature and kerning insertion at the time when the initial node list is built by Lua \TeX 's main control loop. They are part of a temporary trick and will be removed in the near future. For now, you need to 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 now be implemented by overloading the ligature building and kerning functions, i.e. by assigning dummy functions to their associated callbacks.

1.9 `\formatname`

The `\formatname` syntax is identical to `\jobname`. In `ini \TeX` , the expansion is empty. Otherwise, the expansion is the value that `\jobname` had during the `ini \TeX` run that dumped the currently loaded format.

1.10 `\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.
- The '... while end of file ...' error tests are not executed, allowing the expansion to end on a different grouping level or while a conditional is still incomplete.

1.11 Alignments

1.11.1 `\alignmark`

This primitive duplicates the functionality of `#` inside alignment preambles.

1.11.2 `\aligntab`

This primitive duplicates the functionality of `&` inside alignments and preambles.

1.12 Catcode tables

Catcode tables are a new feature that allows you to switch to a predefined catcode regime in a single statement. You can have a practically unlimited number of different tables. 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 tables, and their contents is stored and retrieved from the format file.

1.12.1 `\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 `iniTeX`.

1.12.2 `\initcatcodetable`

`\initcatcodetable` <15-bit number>

The primitive `\initcatcodetable` creates a new table with catcodes identical to those defined by `iniTeX`:

0	<code>\letterbackslash</code>		escape	
5	<code>\letterhat \letterhat M</code>	return	<code>car_ret</code>	(this name may change)
9	<code>\letterhat \letterhat @</code>	null	ignore	
10	<code><space></code>	space	<code>spacer</code>	
11	<code>a - z</code>		letter	
11	<code>A - Z</code>		letter	
12	everything else		other	
14	<code>\letterpercent</code>		comment	
15	<code>\letterhat \letterhat ?</code>	delete	<code>invalid_char</code>	

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.

1.12.3 `\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.

1.13 Suppressing errors

1.13.1 `\suppressfontnotfounderror`

`\suppressfontnotfounderror = 1`

If this new integer parameter is non-zero, then `LuaTeX` will not complain about font metrics that are not found. Instead it will silently skip the font assignment, making the requested csname



for the font `\ifx` equal to `\nullfont`, so that it can be tested against that without bothering the user.

1.13.2 `\suppresslongerror`

```
\suppresslongerror = 1
```

If this new integer parameter is non-zero, then LuaTeX will not complain about `\par` commands encountered in contexts where that is normally prohibited (most prominently in the arguments of non-long macros).

1.13.3 `\suppressifcsnameerror`

```
\suppressifcsnameerror = 1
```

If this new integer parameter is non-zero, then LuaTeX will not complain about non-expandable commands appearing in the middle of a `\ifcsname` expansion. Instead, it will keep getting expanded tokens from the input until it encounters an `\endcsname` command. Use with care! This command is experimental: if the input expansion is unbalanced with respect to `\csname` ...`\endcsname` pairs, the LuaTeX process may hang indefinitely.

1.13.4 `\suppressoutererror`

```
\suppressoutererror = 1
```

If this new integer parameter is non-zero, then LuaTeX will not complain about `\outer` commands encountered in contexts where that is normally prohibited.

1.13.5 `\suppressmathparerror`

The following setting will permit tokens in a math formula:

```
\suppressmathparerror = 1
```

So, the next code is valid then:

```
$ x + 1 =
```

```
a $
```

1.14 `\matheqnogapstep`

By default TeX 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 \TeX .

1.15 `\outputbox`

```
\outputbox = 65535
```

This new 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.

1.16 `\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 5. Here are some more:

```
\bf 14  
\it 18  
\bi 21
```

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 Con \TeX 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).

1.17 `\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.

1.18 `\nohrule` and `\novrule`

Because internally box resources and image resources are now stored as a special kind of rule, we also introduced an empty rule variant. 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).

1.19 `\Uchar`

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



1.20 `\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.

1.21 `\boundary` and `\nboundary`

The `\nboundary` commands 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.

1.22 Debugging

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

1.23 Images and Forms

These two concepts are now core concepts and no longer whatsits. They are in fact now implemented as rules with special properties. Normal rules have subtype 0, saved boxes have subtype 1 and images have subtype 2. This has the positive side effect that whenever we need to take content with dimensions into account, when we look at rule nodes, we automatically also deal with these two types.

The syntax of the `\save...resource` is the same as in pdf_TE_X but you should consider them to be backend specific. This means that a macro package should treat them as such and check for the current output mode if applicable. Here are the equivalents:

```
\saveboxresource           : \pdfxform
\saveimageresource        : \pdfximage
\useboxresource           : \pdfrefxform
\useimageresource        : \pdfrefximage
\lastsavedboxresourceindex : \pdflastxform
\lastsavedimageresourceindex : \pdflastximage
\lastsavedimageresourcepages : \pdflastximagepages
```

Lua_TE_X accepts 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`.

```
\pdfrefximage width 20mm height 10mm depth 5mm \pdflastximage
\pdfrefxform width 20mm height 10mm depth 5mm \pdflastxform
```

1.24 `\outputmode` and `\draftmode`

The `\outputmode` variable tells Lua_T_EX what it has to produce:

```
0 dvi code
1 pdf code
```

The value of the `\draftmode` counter signals the backend if it should output less. The pdf backend accepts a value of 1, while the dvi backend ignores the value.

1.25 File syntax

Lua_T_EX 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.

1.26 Font syntax

Lua_T_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.


1.27 Writing to file

You can now open up to 127 files with `\openout`. When no file is open writes will go to the console and log. As a consequence a system command is no longer possible but one can use `os.execute` to do the same.

1.28 `\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. Below we see the results for four characters separated by a space.





1.29 \letcharcode

This primitive is still experimental but 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).





2 LUA general

2.1 Initialization

2.1.1 L^AT_EX as a LUA interpreter

There are some situations that make Lua_{T_EX} behave like a standalone Lua interpreter:

- if a `--luaonly` option is given on the commandline, or
- if the executable is named `texlua` or `luatexlua`, 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.

Lua_{T_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.

2.1.2 L^AT_EX as a LUA byte compiler

There are two situations that make Lua_{T_EX} behave like the Lua byte compiler:

- if a `--luaonly` option is given on the command line, or
- if the executable is named `texluac`

In this mode, Lua_{T_EX} is exactly like `luac` from the stand alone Lua distribution, except that it does not have the `-l` switch, and that it accepts (but ignores) the `--luaonly` switch.

2.1.3 Other commandline processing

When the Lua_{T_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 in traditional pdf_{T_EX} and Aleph. Some options are accepted but have no consequence. The following command-line options are understood:

<code>--fmt=FORMAT</code>	load the format file <code>FORMAT</code>
<code>--lua=FILE</code>	load and execute a Lua initialization script
<code>--safer</code>	disable easily exploitable Lua commands
<code>--nosocket</code>	disable the Lua socket library
<code>--help</code>	display help and exit
<code>--ini</code>	be <code>iniluatex</code> , for dumping formats
<code>--interaction=STRING</code>	set interaction mode: <code>batchmode</code> , <code>nonstopmode</code> <code>scrollmode</code> or <code>errorstopmode</code>
<code>--halt-on-error</code>	stop processing at the first error



<code>--kpathsea-debug=NUMBER</code>	set path searching debugging flags according to the bits of <code>NUMBER</code>
<code>--progname=STRING</code>	set the program name to <code>STRING</code>
<code>--version</code>	display version and exit
<code>--credits</code>	display credits and exit
<code>--recorder</code>	enable filename recorder
<code>--etex</code>	ignored
<code>--output-comment=STRING</code>	use <code>STRING</code> for dvi file comment instead of date (no effect for pdf)
<code>--output-directory=DIR</code>	use <code>DIR</code> as the directory to write files to
<code>--draftmode</code>	switch on draft mode i.e. generate no output in pdf mode
<code>--output-format=FORMAT</code>	use <code>FORMAT</code> for job output; <code>FORMAT</code> is <code>dvi</code> or <code>pdf</code>
<code>--[no-]shell-escape</code>	disable/enable <code>\write 18{SHELL COMMAND}</code>
<code>--enable-write18</code>	enable <code>\write 18{SHELL COMMAND}</code>
<code>--disable-write18</code>	disable <code>\write 18{SHELL COMMAND}</code>
<code>--shell-restricted</code>	restrict <code>\write 18</code> to a list of commands given in <code>texmf.cnf</code>
<code>--debug-format</code>	enable format debugging
<code>--[no-]file-line-error</code>	disable/enable <code>file:line:error</code> style messages
<code>--[no-]file-line-error-style</code>	aliases of <code>--[no-]file-line-error</code>
<code>--jobname=STRING</code>	set the job name to <code>STRING</code>
<code>--[no-]parse-first-line</code>	ignored
<code>--translate-file=</code>	ignored
<code>--default-translate-file=</code>	ignored
<code>--8bit</code>	ignored
<code>--[no-]mktex=FMT</code>	disable/enable <code>mktexFMT</code> generation with <code>FMT</code> is <code>tex</code> or <code>tfm</code>
<code>--synctex=NUMBER</code>	enable <code>synctex</code>

A note on the creation of the various temporary files and the `\jobname`. 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.
- 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.

The file names for output files that are generated automatically are created by attaching the proper extension (`.log`, `.pdf`, etc.) to the found `\jobname`. These files are created in the directory pointed to by `--output-directory`, or in the current directory, if that switch is not present.

Without the `--lua` option, command line processing works like it does in any other web2c-based typesetting engine, except that Lua_T_EX has a few extra switches.

If the `--lua` option is present, Lua_T_EX will enter an alternative mode of command line processing in comparison to the standard web2c programs.



In this mode, a small series of actions is taken in order. First, it will parse the command line as usual, but it will only interpret a small subset of the options immediately: `--safer`, `--nosocket`, `--[no-]shell-escape`, `--enable-writel8`, `--disable-writel8`, `--shell-restricted`, `--help`, `--version`, and `--credits`.

Now it searches for the requested Lua initialization script. If it cannot be found using the actual name given on the command line, a second attempt is made by prepending the value of the environment variable `LUATEXDIR`, if that variable is defined in the environment.

Then it checks the various safety switches. You can use those to disable some Lua commands that can easily be abused by a malicious document. At the moment, `--safer` `nils` the following functions:

library functions

```
os      execute exec setenv rename remove tmpdir
io      popen output tmpfile
lfs     rmdir mkdir chdir lock touch
```

Furthermore, it disables loading of compiled Lua libraries and it makes `io.open()` fail on files that are opened for anything besides reading.

When Lua_{TEX} starts it set the locale to a neutral value. If for some reason you use `os.locale`, you need to make sure you `nil` it afterwards because otherwise it can interfere with code that for instance generates dates. You can `nil` the locale with

```
os.setlocale(nil,nil)
```

`--nosocket` makes the socket library unavailable, so that Lua cannot use networking.

The switches `--[no-]shell-escape`, `--[enable|disable]-writel8`, and `--shell-restricted` have the same effects as in pdf_{TEX}, and additionally make `io.popen()`, `os.execute`, `os.exec` and `os.spawn` adhere to the requested option.

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, the warning about unrecognized option is suppressed.

Command line processing happens very early on. So early, in fact, that none of _{TEX}'s initializations have taken place yet. For that reason, the tables that deal with typesetting, like `tex`, `token`, `node` and `pdf`, are off-limits during the execution of the startup file (they are `nil`'d). Special care is taken that `texio.write` and `texio.write_nl` function properly, so that you can at least report your actions to the log file when (and if) it eventually becomes opened (note that _{TEX} does not even know its `\jobname` yet at this point). See chapter 7 for more information about the Lua_{TEX}-specific Lua extension tables.

Everything you do in the Lua initialization script will remain visible during the rest of the run, with the exception of the aforementioned `tex`, `token`, `node` and `pdf` tables: those will be initialized to their documented state after the execution of the script. You should not store anything in variables or within tables with these four global names, as they will be overwritten completely.

We recommend you use the startup file only for your own _{TEX}-independent initializations (if you need any), to parse the command line, set values in the `texconfig` table, and register the callbacks you need.



LuaTeX 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).

Unless the `texconfig` table tells LuaTeX not to initialize kpathsea at all (set `texconfig.kpse_init` to `false` for that), LuaTeX acts on some more command line options after the initialization script is finished: in order to initialize the built-in kpathsea library properly, LuaTeX needs to know the correct program name to use, and for that it needs to check `--progrname`, or `--ini` and `--fmt`, if `--progrname` is missing.

2.2 LUA behaviour

Luas `tonumber` function 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`.

Loading dynamic Lua libraries will fail if there are two Lua libraries loaded at the same time (which will typically happen on `win32`, because there is one Lua 5.2 inside LuaTeX, and another will likely be linked to the dll file of the module itself). We plan to fix that later by switching LuaTeX itself to using the dll version of Lua 5.2 inside LuaTeX instead of including a static version in the binary.

LuaTeX is able to use the kpathsea library to find `require()`d modules. For this purpose, `package.searchers[2]` is replaced by a different loader function, that decides at runtime whether to use kpathsea or the built-in core Lua function. It uses kpathsea when that is already initialized at that point in time, otherwise it reverts to using the normal `package.path` loader.

Initialization of kpathsea can happen either implicitly (when LuaTeX starts up and the startup script has not set `texconfig.kpse_init` to `false`), or explicitly by calling the Lua function `kpse.set_program_name()`.

LuaTeX is able to use dynamically loadable Lua libraries, unless `--safer` was given as an option on the command line. For this purpose, `package.searchers[3]` is replaced by a different loader function, that decides at runtime whether to use kpathsea or the built-in core Lua function. It uses kpathsea when that is already initialized at that point in time, otherwise it reverts to using the normal `package.cpath` loader.

This functionality required an extension to kpathsea:

There is a new kpathsea file format: `kpse_clua_format` that searches for files with extension `.dll` and `.so`. The `texmf.cnf` setting for this variable is `CLUAINPUTS`, and by default it has this value:

```
CLUAINPUTS=.:$SELFAUTOLOC/lib/{$progrname,$engine,}/lua//
```

This path is imperfect (it requires a `tds` subtree below the binaries directory), but the architecture has to be in the path somewhere, and the currently simplest way to do that is to search below the binaries directory only. Of course it no big deal to write an alternative loader and use that in a macro package.

One level up (a `lib` directory parallel to `bin`) would have been nicer, but that is not doable because TeXLive uses a `bin/<arch>` structure.



In keeping with the other T_EX-like programs in T_EXLive, the two Lua functions `os.execute` and `io.popen`, as well as the two new functions `os.exec` and `os.spawn` that are explained below, take the value of `shell_escape` and/or `shell_escape_commands` in account. Whenever LuaT_EX is run with the assumed intention to typeset a document (and by that we mean that it is called as `luatex`, as opposed to `texlua`, and that the command line option `--luaonly` was not given), it will only run the four functions above if the matching `texmf.cnf` variable(s) or their `texconfig` (see section 7.15) counterparts allow execution of the requested system command. In ‘script interpreter’ runs of LuaT_EX, these settings have no effect, and all four functions function as normal.

The `f:read("*line")` and `f:lines()` functions from the `io` library have been adjusted so that they are line-ending neutral: any of `LF`, `CR` or `CR+LF` are acceptable line endings.

`luafilesystem` has been extended: there are two extra boolean functions (`lfs.isdir(filename)` and `lfs.isfile(filename)`) and one extra string field in its attributes table (`permissions`). There is an additional function `lfs.shortname()` which takes a file name and returns its short name on `win32` platforms. On other platforms, it just returns the given argument. The file name is not tested for existence. Finally, for non-`win32` platforms only, there is the new function `lfs.readlink()` that takes an existing symbolic link as argument and returns its content. It returns an error on `win32`.

The `string` library has an extra function: `string.explode(s[,m])`. This function 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 T_EX macros.

The `string` library also has six extra iterators that return strings piecemeal:

- `string.utfvalues(s)`: an integer value in the Unicode range
- `string.utfcharacters(s)`: a string with a single utf-8 token in it
- `string.characters(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 utf-16 encoded data into utf-8.

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 `luajit`.

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 `slunicode` 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 will provide some native utf8 support.

The `os` library has a few extra functions and variables:

- `os.selfdir` is a variable that holds the directory path of the actual executable. For example: `\directlua{tex.sprint(os.selfdir)}`.
- `os.exec(commandline)` is a variation on `os.execute`. Here `commandline` can be either a single string or a single table.

If the argument is a table: LuaTeX first checks if there is a value at integer index zero. If there is, this is the command to be executed. Otherwise, it will use the value at integer index one. (if neither are present, nothing at all happens).

The set of consecutive values starting at integer 1 in the table are the arguments that are passed on to the command (the value at index 1 becomes `arg[0]`). The command is searched for in the execution path, so there is normally no need to pass on a fully qualified path name. If the argument is a string, then it is automatically converted into a table by splitting on whitespace. In this case, it is impossible for the command and first argument to differ from each other.

In the string argument format, whitespace can be protected by putting (part of) an argument inside single or double quotes. One layer of quotes is interpreted by LuaTeX, and all occurrences of `\`, `'` or `\\` within the quoted text are unescaped. In the table format, there is no string handling taking place.

This function normally does not return control back to the Lua script: the command will replace the current process. However, it will return the two values `nil` and `'error'` if there was a problem while attempting to execute the command.

On MS Windows, the current process is actually kept in memory until after the execution of the command has finished. This prevents crashes in situations where TeX Lua scripts are run inside integrated TeX environments.

The original reason for this command is that it cleans out the current process before starting the new one, making it especially useful for use in TeX Lua.

- `os.spawn(commandline)` is a returning version of `os.exec`, with otherwise identical calling conventions.

If the command ran ok, then the return value is the exit status of the command. Otherwise, it will return the two values `nil` and `'error'`.

- `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. This function is not available on the SunOS platforms, so do not use this function for portable documents.
- `os.times()` returns the current process times according to the Unix C library function 'times'. This function is not available on the MS Windows and SunOS platforms, so do not use this function for portable documents.
- `os.tmpdir()` creates a directory in the 'current directory' with the name `luatex.XXXXXX`



where the *X*-es are replaced by a unique string. The function also returns this string, so you can `lfs.chdir()` into it, or `nil` if it failed to create the directory. The user is responsible for cleaning up at the end of the run, it does not happen automatically.

- `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` (unknown, but sysv-like), `generic` (unknown).

- `os.version` is planned as a future extension.
- `os.uname()` returns a table with specific operating system information acquired at runtime. The keys in the returned table are all string valued, and their names are: `sysname`, `machine`, `release`, `version`, and `nodename`.

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

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

2.3 LUA modules

The implied use of the built-in Lua modules in this section is deprecated. If you want to use one of these libraries, please start your source file with a proper `require` line. At some point LuaTeX will switch to loading these modules on demand.

Some modules that are normally external to Lua are statically linked in with LuaTeX, because they offer useful functionality:

- `slnunicode`, from the `selene` libraries, <http://luaforge.net/projects/sln>. This library has been slightly extended so that the `unicode.utf8.*` functions also accept the first 256 values of plane 18. This is the range LuaTeX uses for raw binary output, as explained above.
- `luazip`, from the kepler project, <http://www.keplerproject.org/luazip/>.
- `luafilesystem`, also from the kepler project, <http://www.keplerproject.org/luafilesystem/>.
- `lpeg`, by Roberto Ierusalimschy, <http://www.inf.puc-rio.br/~roberto/lpeg/lpeg.html>. 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 utf 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('ää')` results in the message `bad argument #1 to 'R' (range must have two characters)`, since to `lpeg`, `ä` is two ‘characters’ (bytes), so `ää` totals three. In practice this is no real issue.
- `lzlib`, by Tiago Dionizio, <http://luaforge.net/projects/lzlib/>.



- `md5`, by Roberto Ierusalimschy <http://www.inf.puc-rio.br/~roberto/md5/md5-5/md5.html>.
- `luasocket`, by Diego Nehab <http://w3.impa.br/~diego/software/luasocket/>. The `.lua` support modules from `luasocket` are also preloaded inside the executable, there are no external file dependencies.



3 Languages and characters, fonts and glyphs

LuaTeX's internal handling of the characters and glyphs that eventually become typeset is quite different from the way TeX82 handles those same objects. The easiest way to explain the difference is to focus on unrestricted horizontal mode (i.e. paragraphs) and hyphenation first. Later on, it will be easy to deal with the differences that occur in horizontal and math modes.

In TeX82, the characters you type are converted into `char_node` records when they are encountered by the main control loop. TeX attaches and processes the font information while creating those records, so that the resulting 'horizontal list' contains the final forms of ligatures and implicit kerning. This packaging is needed because we may want to get the effective width of for instance a horizontal box.

When it becomes necessary to hyphenate words in a paragraph, TeX converts (one word at time) the `char_node` records into a string array by replacing ligatures with their components and ignoring the kerning. Then it runs the hyphenation algorithm on this string, and converts the hyphenated result back into a 'horizontal list' that is consecutively spliced back into the paragraph stream. Keep in mind that the paragraph may contain unboxed horizontal material, which then already contains ligatures and kerns and the words therein are part of the hyphenation process.

The `char_node` records are somewhat misnamed, as they are glyph positions in specific fonts, and therefore not really 'characters' in the linguistic sense. There is no language information inside the `char_node` records. Instead, language information is passed along using `language_what'sit` records inside the horizontal list.

In LuaTeX, the situation is quite different. The characters you type are always converted into `glyph_node` records with a special subtype to identify them as being intended as linguistic characters. LuaTeX stores the needed language information in those records, but does not do any font-related processing at the time of node creation. It only stores the index of the current font.

When it becomes necessary to typeset a paragraph, LuaTeX first inserts all hyphenation points right into the whole node list. Next, it processes all the font information in the whole list (creating ligatures and adjusting kerning), and finally it adjusts all the subtype identifiers so that the records are 'glyph nodes' from now on.

That was the broad overview. The rest of this chapter will deal with the minutiae of the new process.

3.1 Characters and glyphs

TeX82 (including pdfTeX) differentiates between `char_nodes` and `lig_nodes`. The former are simple items that contained nothing but a 'character' and a 'font' field, and they lived in the same memory as tokens did. The latter also contained a list of components, and a subtype indicating whether this ligature was the result of a word boundary, and it was stored in the same place as other nodes like boxes and kerns and glues.

In LuaTeX, these two types are merged into one, somewhat larger structure called a `glyph_node`. Besides having the old character, font, and component fields, and the new special fields like 'attr' (see section 6.1.2.12), these nodes also contain:



- A subtype, split into four main types:
 - `character`, for characters to be hyphenated: the lowest bit (bit 0) is set to 1.
 - `glyph`, for specific font glyphs: the lowest bit (bit 0) is not set.
 - `ligature`, for ligatures (bit 1 is set)
 - `ghost`, for ‘ghost objects’ (bit 2 is set)

The latter two make further use of two extra fields (bits 3 and 4):

- `left`, for ligatures created from a left word boundary and for ghosts created from `\left-ghost`
- `right`, for ligatures created from a right word boundary and for ghosts created from `\rightghost`

For ligatures, both bits can be set at the same time (in case of a single-glyph word).

- `glyph_nodes` of type ‘character’ also contain language data, split into four items that were current when the node was created: the `\setlanguage` (15 bits), `\lefthyphenmin` (8 bits), `\righthyphenmin` (8 bits), and `\uchyph` (1 bit).

Incidentally, Lua \TeX allows 16383 separate languages, and words can be 256 characters long. The language is stored with each character. You can set `\firstvalidlanguage` to for instance 1 and make thereby language 0 an ignored hyphenation language.

The new primitive `\hyphenationmin` can be used to signal the minimal length of a word. This value stored with the (current) language.

Because the `\uchyph` value is saved in the actual nodes, its handling is subtly different from \TeX 82: changes to `\uchyph` become effective immediately, not at the end of the current partial paragraph.

Typeset boxes now always have their language information embedded in the nodes themselves, so there is no longer a possible dependency on the surrounding language settings. In \TeX 82, a mid-paragraph statement like `\unhbox0` would process the box using the current paragraph language unless there was a `\setlanguage` issued inside the box. In Lua \TeX , all language variables are already frozen.

In traditional \TeX the process of hyphenation is driven by so called lccodes. In Lua \TeX we made this dependency less strong. There are several strategies possible. When you do nothing, the currently used lccodes are used, when loading patterns, setting exceptions or hyphenating a list.

When you set `\savingshyphcodes` to a value larger than zero the current set of lccodes will be saved with the language. In that case changing a lccode afterwards has no effect. However, you can adapt the set with:

```
\hjcode`a=`a
```

This change is global which makes sense if you keep in mind that the moment that hyphenation happens is (normally) when the paragraph or a horizontal box is constructed. When `\savingshyphcodes` was zero when the language got initialized you start out with nothing, otherwise you already have a set.

Carrying all this information with each glyph would give too much overhead and also make the definition more complex. A solution with hj codesets was considered but rejected because in practice the current approach is sufficient and it would not be compatible anyway.



Beware: the values are always saved in the format, independent of the setting of `\savingshyphcodes` at the moment the format is dumped.

3.2 The main control loop

In Lua_T_EX's main loop, almost all input characters that are to be typeset are converted into `glyph` node records with subtype 'character', but there are a few exceptions.

First, the `\accent` primitive creates nodes with subtype 'glyph' instead of 'character': one for the actual accent and one for the accentee. The primary reason for this is that `\accent` in T_EX82 is explicitly dependent on the current font encoding, so it would not make much sense to attach a new meaning to the primitive's name, as that would invalidate many old documents and macro packages. A secondary reason is that in T_EX82, `\accent` prohibits hyphenation of the current word. Since in Lua_T_EX hyphenation only takes place on 'character' nodes, it is possible to achieve the same effect.

This change of meaning did happen with `\char`, that now generates 'glyph' nodes with a character subtype. In traditional T_EX there was a strong relationship between the 8-bit input encoding, hyphenation and glyphs taken from a font. In Lua_T_EX we have utf input, and in most cases this maps directly to a character in a font, apart from glyph replacement in the font engine. If you want to access arbitrary glyphs in a font directly you can always use Lua to do so, because fonts are available as Lua table.

Second, all the results of processing in math mode eventually become nodes with 'glyph' subtypes.

Third, the Aleph-derived commands `\leftghost` and `\rightghost` create nodes of a third subtype: 'ghost'. These nodes are ignored completely by all further processing until the stage where inter-glyph kerning is added.

Fourth, automatic discretionary are handled differently. T_EX82 inserts an empty discretionary after sensing an input character that matches the `\hyphenchar` in the current font. This test is wrong, in our opinion: whether or not hyphenation takes place should not depend on the current font, it is a language property.

In Lua_T_EX, it works like this: if Lua_T_EX senses a string of input characters that matches the value of the new integer parameter `\exhyphenchar`, it will insert an explicit discretionary after that series of nodes. Initex sets the `\exhyphenchar=-1`. Incidentally, this is a global parameter instead of a language-specific one because it may be useful to change the value depending on the document structure instead of the text language.

The insertion of discretionary after a sequence of explicit hyphens happens at the same time as the other hyphenation processing, *not* inside the main control loop.

The only use Lua_T_EX has for `\hyphenchar` is at the check whether a word should be considered for hyphenation at all. If the `\hyphenchar` of the font attached to the first character node in a word is negative, then hyphenation of that word is abandoned immediately. **This behaviour is added for backward compatibility only, and the use of `\hyphenchar=-1` as a means of preventing hyphenation should not be used in new Lua_T_EX documents.**

Fifth, `\setlanguage` no longer creates whatsits. The meaning of `\setlanguage` is changed so that it is now an integer parameter like all others. That integer parameter is used in `\glyph_node`



creation to add language information to the glyph nodes. In conjunction, the `\language` primitive is extended so that it always also updates the value of `\setlanguage`.

Sixth, the `\noboundary` command (this command prohibits word boundary processing where that would normally take place) now does create whatsits. These whatsits are needed because the exact place of the `\noboundary` command in the input stream has to be retained until after the ligature and font processing stages.

Finally, there is no longer a `main_loop` label in the code. Remember that T_EX82 did quite a lot of processing while adding `char_nodes` to the horizontal list? For speed reasons, it handled that processing code outside of the ‘main control’ loop, and only the first character of any ‘word’ was handled by that ‘main control’ loop. In LuaT_EX, there is no longer a need for that (all hard work is done later), and the (now very small) bits of character-handling code have been moved back inline. When `\tracingcommands` is on, this is visible because the full word is reported, instead of just the initial character.

3.3 Loading patterns and exceptions

The hyphenation algorithm in LuaT_EX is quite different from the one in T_EX82, although it uses essentially the same user input.

After expansion, the argument for `\patterns` has to be proper utf8 with individual patterns separated by spaces, no `\char` or `\chardef` commands are allowed. The current implementation is even more strict, and will reject all non-Unicode characters, but that will be changed in the future. For now, the generated errors are a valuable tool in discovering font-encoding specific pattern files.

Likewise, the expanded argument for `\hyphenation` also has to be proper utf8, but here a tiny little bit of extra syntax is provided:

1. Three sets of arguments in curly braces (`{ } { } { }`) indicates a desired complex discretionary, with arguments as in `\discretionary`’s command in normal document input.
2. A `-` indicates a desired simple discretionary, cf. `\-` and `\discretionary{-}{ } { }` in normal document input.
3. Internal command names are ignored. This rule is provided especially for `\discretionary`, but it also helps to deal with `\relax` commands that may sneak in.
4. An `=` indicates a (non-discretionary) hyphen in the document input.

The expanded argument is first converted back to a space-separated string while dropping the internal command names. This string is then converted into a dictionary by a routine that creates key-value pairs by converting the other listed items. It is important to note that the keys in an exception dictionary can always be generated from the values. Here are a few examples:

value	implied key (input)	effect
<code>ta-ble</code>	table	<code>ta\ble (= ta\discretionary{-}{ } { }ble)</code>
<code>ba{k-}{ } {c}ken</code>	backen	<code>ba\discretionary{k-}{ } {c}ken</code>

The resultant patterns and exception dictionary will be stored under the language code that is the present value of `\language`.



In the last line of the table, you see there is no `\discretionary` command in the value: the command is optional in the \TeX -based input syntax. The underlying reason for that is that it is conceivable that a whole dictionary of words is stored as a plain text file and loaded into Lua \TeX using one of the functions in the Lua `lang` library. This loading method is quite a bit faster than going through the \TeX language primitives, but some (most?) of that speed gain would be lost if it had to interpret command sequences while doing so.

It is possible to specify extra hyphenation points in compound words by using `{-}{}{-}` for the explicit hyphen character (replace `-` by the actual explicit hyphen character if needed). For example, this matches the word ‘multi-word-boundaries’ and allows an extra break inbetween ‘boun’ and ‘daries’:

```
\hyphenation{multi{-}{}{-}word{-}{}{-}boun-daries}
```

The motivation behind the ε - \TeX extension `\savingshyphcodes` was that hyphenation heavily depended on font encodings. This is no longer true in Lua \TeX , and the corresponding primitive is ignored pending complete removal. The future semantics of `\uppercase` and `\lowercase` are still under consideration, no changes have taken place yet.

3.4 Applying hyphenation

The internal structures Lua \TeX uses for the insertion of discretionaries in words is very different from the ones in \TeX 82, and that means there are some noticeable differences in handling as well.

First and foremost, there is no ‘compressed trie’ involved in hyphenation. The algorithm still reads patgen-generated pattern files, but Lua \TeX uses a finite state hash to match the patterns against the word to be hyphenated. This algorithm is based on the ‘libhunj’ library used by OpenOffice, which in turn is inspired by \TeX . The memory allocation for this new implementation is completely dynamic, so the web2c setting for `trie_size` is ignored.

Differences between Lua \TeX and \TeX 82 that are a direct result of that:

- Lua \TeX happily hyphenates the full Unicode character range.
- Pattern and exception dictionary size is limited by the available memory only, all allocations are done dynamically. The trie-related settings in `texmf.cnf` are ignored.
- Because there is no ‘trie preparation’ stage, language patterns never become frozen. This means that the primitive `\patterns` (and its Lua counterpart `lang.patterns`) can be used at any time, not only in `ini \TeX` .
- Only the string representation of `\patterns` and `\hyphenation` is stored in the format file. At format load time, they are simply re-evaluated. It follows that there is no real reason to preload languages in the format file. In fact, it is usually not a good idea to do so. It is much smarter to load patterns no sooner than the first time they are actually needed.
- Lua \TeX uses the language-specific variables `\prehyphenchar` and `\posthyphenchar` in the creation of implicit discretionaries, instead of \TeX 82’s `\hyphenchar`, and the values of the language-specific variables `\preexhyphenchar` and `\postexhyphenchar` for explicit discretionaries (instead of \TeX 82’s empty discretionary).
- The value of the two counters related to hyphenation, `\hyphenpenalty` and `\exhyphenpenalty`, are now stored in the discretionary nodes. This permits a local overload for explicit



`\discretionary` commands. The value current when the hyphenation pass is applied is used. When no callbacks are used this is compatible with traditional \TeX . When you apply the Lua `lang.hyphenate` function the current values are used.

Because we store penalties in the disc node the `\discretionary` command has been extended to accept an optional penalty specification, so you can do the following:

```
\hsizelmm
1:foo{\hyphenpenalty 10000\discretionary{}{}{}}bar\par
2:foo\discretionary penalty 10000 {}{}{}}bar\par
3:foo\discretionary{}{}{}}bar\par
```

This results in:

```
1:foobar
2:foobar
3:foo
bar
```

Inserted characters and ligatures inherit their attributes from the nearest glyph node item (usually the preceding one, but the following one for the items inserted at the left-hand side of a word).

Word boundaries are no longer implied by font switches, but by language switches. One word can have two separate fonts and still be hyphenated correctly (but it can not have two different languages, the `\setlanguage` command forces a word boundary).

All languages start out with `\prehyphenchar=\-`, `\posthyphenchar=0`, `\preexhyphenchar=0` and `\postexhyphenchar=0`. When you assign the values of one of these four parameters, you are actually changing the settings for the current `\language`, this behaviour is compatible with `\patterns` and `\hyphenation`.

Lua \TeX also hyphenates the first word in a paragraph. Words can be up to 256 characters long (up from 64 in \TeX 82). Longer words generate an error right now, but eventually either the limitation will be removed or perhaps it will become possible to silently ignore the excess characters (this is what happens in \TeX 82, but there the behaviour cannot be controlled).

If you are using the Lua function `lang.hyphenate`, you should be aware that this function expects to receive a list of ‘character’ nodes. It will not operate properly in the presence of ‘glyph’, ‘ligature’, or ‘ghost’ nodes, nor does it know how to deal with kerning. In the near future, it will be able to skip over ‘ghost’ nodes, and we may add a less fuzzy function you can call as well.

The hyphenation exception dictionary is maintained as key-value hash, and that is also dynamic, so the `hyph_size` setting is not used either.

3.5 Applying ligatures and kerning

After all possible hyphenation points have been inserted in the list, Lua \TeX will process the list to convert the ‘character’ nodes into ‘glyph’ and ‘ligature’ nodes. This is actually done in two stages: first all ligatures are processed, then all kerning information is applied to the result list.



But those two stages are somewhat dependent on each other: If the used font makes it possible to do so, the ligaturing stage adds virtual ‘character’ nodes to the word boundaries in the list. While doing so, it removes and interprets `noboundary` nodes. The kerning stage deletes those word boundary items after it is done with them, and it does the same for ‘ghost’ nodes. Finally, at the end of the kerning stage, all remaining ‘character’ nodes are converted to ‘glyph’ nodes.

This work separation is worth mentioning because, if you overrule from Lua only one of the two callbacks related to font handling, then you have to make sure you perform the tasks normally done by LuaTeX itself in order to make sure that the other, non-overruled, routine continues to function properly.

Work in this area is not yet complete, but most of the possible cases are handled by our rewritten ligaturing engine. We are working hard to make sure all of the possible inputs will become supported soon.

For example, take the word `office`, hyphenated `of-fice`, using a ‘normal’ font with all the `f-f` and `f-i` type ligatures:

```
Initial:           {o}{f}{f}{i}{c}{e}
After hyphenation: {o}{f}{{-}, {}, {}}{f}{i}{c}{e}
First ligature stage: {o}{{f-}, {f}, {<ff>}}{i}{c}{e}
Final result:       {o}{{f-}, {<fi>}, {<ffi>}}{c}{e}
```

That’s bad enough, but let us assume that there is also a hyphenation point between the `f` and the `i`, to create `of-f-ice`. Then the final result should be:

```
{o}{{f-},
  {{f-},
   {i},
   {<fi>}},
  {{<ff>-},
   {i},
   {<ffi>}}}{c}{e}
```

with discretionaries in the post-break text as well as in the replacement text of the top-level discretionary that resulted from the first hyphenation point.

Here is that nested solution again, in a different representation:

	pre	post	replace
topdisc	<code>f⁻¹</code>	<code>sub1</code>	<code>sub2</code>
sub1	<code>f⁻²</code>	<code>i³</code>	<code><fi>⁴</code>
sub2	<code><ff>⁻⁵</code>	<code>i⁶</code>	<code><ffi>⁷</code>

When line breaking is choosing its breakpoints, the following fields will eventually be selected:

```
of-f-ice  f-1
          f-2
          i3
of-fice   f-1
          <fi>4
```



```

off-ice  <ff>-5
         i6
office   <ffi>7

```

The current solution in Lua_T_EX is not able to handle nested discretionaries, but it is in fact smart enough to handle this fictional `of-f-ice` example. It does so by combining two sequential discretionary nodes as if they were a single object (where the second discretionary node is treated as an extension of the first node).

One can observe that the `of-f-ice` and `off-ice` cases both end with the same actual post replacement list (`i`), and that this would be the case even if that `i` was the first item of a potential following ligature like `ic`. This allows Lua_T_EX to do away with one of the fields, and thus make the whole stuff fit into just two discretionary nodes.

The mapping of the seven list fields to the six fields in this discretionary node pair is as follows:

field	description
<code>discl.pre</code>	<code>f</code> ⁻¹
<code>discl.post</code>	<code><fi></code> ⁴
<code>discl.replace</code>	<code><ffi></code> ⁷
<code>disc2.pre</code>	<code>f</code> ⁻²
<code>disc2.post</code>	<code>i</code> ^{3,6}
<code>disc2.replace</code>	<code><ff>-</code> ⁵

What is actually generated after ligaturing has been applied is therefore:

```

{0}{{f-},
     {<fi>},
     {<ffi>}}
{{f-},
 {i},
 {<ff>-}}{c}{e}

```

The two discretionaries have different subtypes from a discretionary appearing on its own: the first has subtype 4, and the second has subtype 5. The need for these special subtypes stems from the fact that not all of the fields appear in their ‘normal’ location. The second discretionary especially looks odd, with things like the `<ff>-` appearing in `disc2.replace`. The fact that some of the fields have different meanings (and different processing code internally) is what makes it necessary to have different subtypes: this enables Lua_T_EX to distinguish this sequence of two joined discretionary nodes from the case of two standalone discretionaries appearing in a row.

Of course there is still that relationship with fonts: ligatures can be implemented by mapping a sequence of glyphs onto one glyph, but also by selective replacement and kerning. This means that the above examples are just representing the traditional approach.

3.6 Breaking paragraphs into lines

This code is still almost unchanged, but because of the above-mentioned changes with respect to discretionaries and ligatures, line breaking will potentially be different from traditional _T_EX.



The actual line breaking code is still based on the T_EX82 algorithms, and it does not expect there to be discretionary nodes inside of discretionary nodes.

But that situation is now fairly common in LuaT_EX, due to the changes to the ligaturing mechanism. And also, the LuaT_EX discretionary nodes are implemented slightly different from the T_EX82 nodes: the `no_break` text is now embedded inside the disc node, where previously these nodes kept their place in the horizontal list (the discretionary node contained a counter indicating how many nodes to skip).

The combined effect of these two differences is that LuaT_EX does not always use all of the potential breakpoints in a paragraph, especially when fonts with many ligatures are used.





4 Font structure

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 returned by the `define_font` callback, or if they result from the normal tfm/vf reading routines if there is no `define_font` callback defined.

The column ‘from vf’ means that this key will be created by the `font.read_vf()` routine, ‘from tfm’ means that the key will be created by the `font.read_tfm()` routine, and ‘used’ means whether or not the LuaT_EX engine itself will do something with the key.

The top-level keys in the table are as follows:

key	from vf	from tfm	used	value type	description
name	yes	yes	yes	string	metric (file) name
area	no	yes	yes	string	(directory) location, typically empty
used	no	yes	yes	boolean	used already? (initial: false)
characters	yes	yes	yes	table	the defined glyphs of this font
checksum	yes	yes	no	number	default: 0
designsize	no	yes	yes	number	expected size (default: 655360 == 10pt)
direction	no	yes	yes	number	default: 0 (TLT)
encodingbytes	no	no	yes	number	default: depends on <code>format</code>
encodingname	no	no	yes	string	encoding name
fonts	yes	no	yes	table	locally used fonts
psname	no	no	yes	string	actual (PostScript) name (this is the PS fontname in the incoming font source, also used as fontname identifier in the pdf output)
fullname	no	no	yes	string	output font name, used as a fall-back in the pdf output if the psname is not set
header	yes	no	no	string	header comments, if any
hyphenchar	no	no	yes	number	default: TeX’s <code>\hyphenchar</code>
parameters	no	yes	yes	hash	default: 7 parameters, all zero
size	no	yes	yes	number	loaded (at) size. (default: same as designsize)
skewchar	no	no	yes	number	default: TeX’s <code>\skewchar</code>
type	yes	no	yes	string	basic type of this font
format	no	no	yes	string	disk format type
embedding	no	no	yes	string	pdf inclusion
filename	no	no	yes	string	disk file name
tounicode	no	yes	yes	number	if 1, LuaT _E X assumes per-glyph tounicode entries are present in the font
stretch	no	no	yes	number	the ‘stretch’ value from <code>\expandglyphsinfont</code>



shrink	no	no	yes	number	the ‘shrink’ value from <code>\expandglyphsinfont</code>
step	no	no	yes	number	the ‘step’ value from <code>\expandglyphsinfont</code>
auto_expand	no	no	yes	boolean	the ‘autoexpand’ keyword from <code>\expandglyphsinfont</code>
expansion_factor	no	no	no	number	the actual expansion factor of an expanded font
attributes	no	no	yes	string	the <code>\pdffontattr</code>
cache	no	no	yes	string	this key controls caching of the Lua table on the <code>tex</code> end. yes : use a reference to the table that is passed to LuaTeX (this is the default). no : don’t store the table reference, don’t cache any Lua data for this font. renew : don’t store the table reference, but save a reference to the table that is created at the first access to one of its fields in <code>font.fonts</code> . Note: the saved reference is thread-local, so be careful when you are using coroutines: an error will be thrown if the table has been cached in one thread, but you reference it from another thread (\approx coroutine)
nomath	no	no	yes	boolean	this key allows a minor speedup for text fonts. if it is present and true, then LuaTeX will not check the character entries for math-specific keys.
slant	no	no	yes	number	This has the same semantics as the <code>SlantFont</code> operator in font map files.
extent	no	no	yes	number	This has the same semantics as the <code>ExtendFont</code> operator in font map files.

The key `name` is always required. The keys `stretch`, `shrink`, `step` and optionally `auto_expand` only have meaning when used together: they can be used to replace a post-loading `\expandglyphsinfont` command. The `expansion_factor` is value that can be present inside a font in `font.fonts`. It is the actual expansion factor (a value between `-shrink` and `stretch`, with step `step`) of a font that was automatically generated by the font expansion algorithm. The key `attributes` can be used to replace `\pdffontattr`. The key `used` is set by the engine when a font is actively in use, this makes sure that the font’s definition is written to the output file (dvi or pdf). The tfm reader sets it to false. The `direction` is a number signalling the ‘normal’ direction for



this font. There are sixteen possibilities:

number	meaning	number	meaning
0	LT	8	TT
1	LL	9	TL
2	LB	10	TB
3	LR	11	TR
4	RT	12	BT
5	RL	13	BL
6	RB	14	BB
7	RR	15	BR

These are Omega-style direction abbreviations: the first character indicates the ‘first’ edge of the character glyphs (the edge that is seen first in the writing direction), the second the ‘top’ side.

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	internal remapped number
slant	1
space	2
space_stretch	3
space_shrink	4
x_height	5
quad	6
extra_space	7

The keys `type`, `format`, `embedding`, `fullname` and `filename` are used to embed OpenType fonts in the result pdf.

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.

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!).

Other index keys are ignored.

Each character hash itself is a hash. For example, here is the character ‘f’ (decimal 102) in the font `cmr10` at 10 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
    }
  }
}
}

```

The following top-level keys can be present inside a character hash:

key	from vf	from tfm	used	value type	description
width	yes	yes	yes	number	character's width, in sp (default 0)
height	no	yes	yes	number	character's height, in sp (default 0)
depth	no	yes	yes	number	character's depth, in sp (default 0)
italic	no	yes	yes	number	character's italic correction, in sp (default zero)
top_accent	no	no	maybe	number	character's top accent alignment place, in sp (default zero)
bot_accent	no	no	maybe	number	character's bottom accent alignment place, in sp (default zero)
left_protruding	no	no	maybe	number	character's <code>\lrcode</code>
right_protruding	no	no	maybe	number	character's <code>\rrcode</code>
expansion_factor	no	no	maybe	number	character's <code>\efcode</code>
tounicode	no	no	maybe	string	character's Unicode equivalent(s), in utf-16BE hexadecimal format
next	no	yes	yes	number	the 'next larger' character index
extensible	no	yes	yes	table	the constituent parts of an extensible recipe



<code>vert_variants</code>	no	no	yes	table	constituent parts of a vertical variant set
<code>horiz_variants</code>	no	no	yes	table	constituent parts of a horizontal variant set
<code>kerns</code>	no	yes	yes	table	kerning information
<code>ligatures</code>	no	yes	yes	table	ligaturing information
<code>commands</code>	yes	no	yes	array	virtual font commands
<code>name</code>	no	no	no	string	the character (PostScript) name
<code>index</code>	no	no	yes	number	the (OpenType or TrueType) font glyph index
<code>used</code>	no	yes	yes	boolean	typeset already (default: false)?
<code>mathkern</code>	no	no	yes	table	math cut-in specifications

The values of `top_accent`, `bot_accent` and `mathkern` are used only for math accent and superscript placement, see the math chapter 59 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`.

The usage of `tounicode` is this: if this font specifies a `tounicode=1` at the top level, then LuaTeX will construct a `/ToUnicode` entry for the pdf font (or font subset) based on the character-level `tounicode` strings, where they are available. If a character does not have a sensible Unicode equivalent, do not provide a string either (no empty strings).

If the font level `tounicode` is not set, then LuaTeX will build up `/ToUnicode` based on the TeX code points you used, and any character-level `tounicodes` will be ignored. The string format is exactly the format that is expected by Adobe CMap files (utf-16BE in hexadecimal encoding), minus the enclosing angle brackets. Small example: the `tounicode` for a `fi` ligature would be `00660069`. When you pass a number the conversion will be done for you.

The presence of `extensible` will overrule `next`, if that is also present. It in turn can be overruled by `vert_variants`.

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	Maximum overlap at the starting side (in scaled points).



end number Maximum overlap at the ending side (in scaled points).
 advance number The total advance width of this item (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 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
<code>type</code>	number	the type of this ligature command, default 0
<code>char</code>	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 T_EX. When T_EX 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.

The `commands` array is explained below.

4.1 Real fonts

Whether or not a T_EX font is a ‘real’ font that should be written to the pdf document is decided by the `type` value in the top-level font structure. If the value is `real`, then this is a proper font, and the inclusion mechanism will attempt to add the needed font object definitions to the pdf.

Values for `type`:

value	description
<code>real</code>	this is a base font
<code>virtual</code>	this is a virtual font



The actions to be taken depend on a number of different variables:

- Whether the used font fits in an 8-bit encoding scheme or not.
- The type of the disk font file.
- The level of embedding requested.

A font that uses anything other than an 8-bit encoding vector has to be written to the pdf in a different way.

The rule is: if the font table has `encodingbytes` set to 2, then this is a wide font, in all other cases it isn't. The value 2 is the default for OpenType and TrueType fonts loaded via Lua. For Type1 fonts, you have to set `encodingbytes` to 2 explicitly. For pk bitmap fonts, wide font encoding is not supported at all.

If no special care is needed, LuaTeX currently falls back to the mapfile-based solution used by pdfTeX and dvips. This behaviour will be removed in the future, when the existing code becomes integrated in the new subsystem.

But if this is a 'wide' font, then the new subsystem kicks in, and some extra fields have to be present in the font structure. In this case, LuaTeX does not use a map file at all.

The extra fields are: `format`, `embedding`, `fullname`, `cidinfo` (as explained above), `filename`, and the `index` key in the separate characters.

Values for `format` are:

value	description
<code>type1</code>	this is a PostScript Type1 font
<code>type3</code>	this is a bitmapped (pk) font
<code>truetype</code>	this is a TrueType or TrueType-based OpenType font
<code>opentype</code>	this is a PostScript-based OpenType font

`type3` fonts are provided for backward compatibility only, and do not support the new wide encoding options.

Values for `embedding` are:

value	description
<code>no</code>	don't embed the font at all
<code>subset</code>	include and attempt to subset the font
<code>full</code>	include this font in its entirety

It is not possible to artificially modify the transformation matrix for the font at the moment.

The other fields are used as follows: The `fullname` will be the PostScript/pdf font name. The `cidinfo` will be used as the character set (the CID `/Ordering` and `/Registry` keys). The `filename` points to the actual font file. If you include the full path in the `filename` or if the file is in the local directory, LuaTeX will run a little bit more efficient because it will not have to re-run the `find_XXX_file` callback in that case.

Be careful: when mixing old and new fonts in one document, it is possible to create PostScript name clashes that can result in printing errors. When this happens, you have to change the `fullname` of the font.



Typeset strings are written out in a wide format using 2 bytes per glyph, using the `index` key in the character information as value. The overall effect is like having an encoding based on numbers instead of traditional (PostScript) name-based reencoding. The way to get the correct `index` numbers for Type1 fonts is by loading the font via `fontloader.open`; use the table indices as `index` fields.

This type of reencoding means that there is no longer a clear connection between the text in your input file and the strings in the output pdf file. Dealing with this is high on the agenda.

4.2 Virtual fonts

You have to take the following steps if you want LuaTeX to treat the returned table from `define_font` as a virtual font:

- Set the top-level key `type` to `virtual`.
- Make sure there is at least one valid entry in `fonts` (see below).
- Give a `commands` array to every character (see below).

The presence of the toplevel `type` key with the specific value `virtual` will trigger handling of the rest of the special virtual font fields in the table, but the mere existence of 'type' is enough to prevent LuaTeX from looking for a virtual font on its own.

Therefore, this also works 'in reverse': if you are absolutely certain that a font is not a virtual font, assigning the value `base` or `real` to `type` will inhibit LuaTeX from looking for a virtual font file, thereby saving you a disk search.

The `fonts` is another Lua array. 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, you can use the `font.nextid()` function which returns the index of the next to be defined font which is probably the currently defined one.

An example makes this easy to understand

```
fonts = {
  { name = 'ptmr8a', size = 655360 },
  { name = 'psyr', size = 600000 },
  { id = 38 }
}
```

says that the first referenced font (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 previously defined font that is known to LuaTeX 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 name	arguments	arg type	description
font	1	number	select a new font from the local <code>fonts</code> 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	number	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 <code>\special</code> command
lua	1	string	execute a Lua script (at <code>\latelua</code> time)
image	1	image	output an image (the argument can be either an <code><image></code> variable or an <code>image_spec</code> 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 (normally one could use `font.nextid` but when more complex fonts are built in the meantime other instances could have been loaded.

Here is a rather elaborate glyph commands example:

```
...
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')
  { 'pop' },           -- go all the way back
  { 'down', -200000 }, -- move upwards by about 3pt
  { 'special', 'pdf: 1 0 0 rg' } -- switch to red color
  { 'rule', 500000, 20000 } -- draw a bar
  { 'special', 'pdf: 0 g' } -- back to black
}
...
```

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 create 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.



4.2.1 Artificial fonts

Even in a ‘real’ font, there can be virtual characters. When LuaT_EX encounters a `commands` field inside a character when it becomes time to typeset the character, it will interpret the commands, just like for a true virtual character. In this case, if you have created no ‘fonts’ array, then the default (and only) ‘base’ font is taken to be the current font itself. In practice, this means that you can create virtual duplicates of existing characters which is useful if you want to create composite characters.

Note: this feature does *not* work the other way around. There can not be ‘real’ characters in a virtual font! You cannot use this technique for font re-encoding either; you need a truly virtual font for that (because characters that are already present cannot be altered).

4.2.2 Example virtual font

Finally, here is a plain T_EX input file with a virtual font demonstration:

```
\directlua {
  callback.register('define_font',
    function (name,size)
      if name == 'cmr10-red' then
        f = font.read_tfm('cmr10',size)
        f.name = 'cmr10-red'
        f.type = 'virtual'
        f.fonts = {{ name = 'cmr10', size = size }}
        for i,v in pairs(f.characters) do
          if (string.char(i)):find('[tacohanshartmut]') then
            v.commands = {
              {'special','pdf: 1 0 0 rg'},
              {'char',i},
              {'special','pdf: 0 g'},
            }
          else
            v.commands = {'char',i}
          end
        end
      end
      f = font.read_tfm(name,size)
    end
  )
}

\font\myfont = cmr10-red at 10pt \myfont This is a line of text \par
\font\myfontx= cmr10 at 10pt \myfontx Here is another line of text \par
```



5 Math

The handling of mathematics in Lua \TeX differs quite a bit from how \TeX 82 (and therefore pdf \TeX) handles math. First, Lua \TeX adds primitives and extends some others so that Unicode input can be used easily. Second, all of \TeX 82'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 in the past that are now added to the engine.

5.1 The current math style

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, Lua \TeX adds the new primitive: `\mathstyle`. This is a 'convert command' like e.g. `\romannumeral`: its value can only be read, not set.

5.1.1 `\mathstyle`

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
```

5.1.2 `\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, 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.

5.2 Unicode math characters

Character handling is now extended up to the full Unicode range (the `\U` prefix), which is compatible with X \TeX .



The math primitives from T_EX 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 LuaT_EX primitives read more than one separate value. This is shown in the tables below by a plus sign in the second column.

The input for such primitives would look like this:

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

Altered T_EX82 primitives:

primitive	value range (in hex)
<code>\mathcode</code>	0-10FFFF = 0-8000
<code>\delcode</code>	0-10FFFF = 0-FFFFFF

Unaltered:

primitive	value range (in hex)
<code>\mathchardef</code>	0-8000
<code>\mathchar</code>	0-7FFF
<code>\mathaccent</code>	0-7FFF
<code>\delimiter</code>	0-7FFFFFFF
<code>\radical</code>	0-7FFFFFFF

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

New primitives that are compatible with X_YT_EX:

primitive	value range (in hex)
<code>\Umathchardef</code>	0+0+0-7+FF+10FFFF ¹
<code>\Umathcharnumdef</code> ⁵	-80000000-7FFFFFFF ³
<code>\Umathcode</code>	0-10FFFF = 0+0+0-7+FF+10FFFF ¹
<code>\Udelcode</code>	0-10FFFF = 0+0-FF+10FFFF ²
<code>\Umathchar</code>	0+0+0-7+FF+10FFFF
<code>\Umathaccent</code>	0+0+0-7+FF+10FFFF ^{2,4}
<code>\Udelimiter</code>	0+0+0-7+FF+10FFFF ²
<code>\Uradical</code>	0+0-FF+10FFFF ²
<code>\Umathcharnum</code>	-80000000-7FFFFFFF ³
<code>\Umathcodenum</code>	0-10FFFF = -80000000-7FFFFFFF ³
<code>\Udelcodenum</code>	0-10FFFF = -80000000-7FFFFFFF ³

Specifications typically look like:

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

Note 1: 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 `\Umathoperatorsizes` parameter (more information can be found in a following section).



Note 2: For these three 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.

Note 3: The `\Umathaccent` command accepts optional keywords to control various details regarding math accents. See section 5.8 below for details.

New primitives that exist in LuaTeX only (all of these will be explained in following sections):

primitive	value range (in hex)
<code>\Uroot</code>	0+0-FF+10FFFF ²
<code>\Uoverdelimiter</code>	0+0-FF+10FFFF ²
<code>\Uunderdelimiter</code>	0+0-FF+10FFFF ²
<code>\Udelimiterover</code>	0+0-FF+10FFFF ²
<code>\Udelimiterunder</code>	0+0-FF+10FFFF ²

5.3 Cramped math styles

LuaTeX 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 taken 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 LuaTeX 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. If we force styles in the script using `\scriptstyle` and `\crampedscriptstyle` we get this:



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:

default	$b_{x=xx}^{x=xx}$	$=x$	x
script	$b_x^{x=xx}$	$=x$	x
crampedscript	$b_{x=xx}^{x=xx}$	$=x$	x

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:

default	$b_x^{x=xx}$	$=x$	x
script	$b_x^{x=xx}$	$=x$	x
crampedscript	$b_x^{x=xx}$	$=x$	x

5.4 Math parameter settings

In LuaT_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 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>\dotswithdelims</code>
<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>\Umathsupshiftp</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.



5.5 Skips around display math

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 later makes it sometimes hard to fully control spacing. Therefore LuaTeX comes with a new directive: `\mathdisplayskipmode`. The following values apply:

- 0 normal TeX behaviour: always above, only below when larger than zero
- 1 always
- 2 only when not zero
- 3 never, not even when not zero

5.6 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, LuaTeX 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 using in the TeXbook. Assignments to `\textfont` set the values for the cramped and uncramped display and text styles. Use `\scriptfont` for the script styles, and `\scriptscriptfont` for the scriptscript styles (totalling 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, LuaTeX also looks at the ‘space’ font dimension parameter. For math fonts, this should be set to zero.

variable	style	default value opentype	default value tfm
<code>\Umathaxis</code>	-	AxisHeight	axis_height
<code>\Umathoperator</code>	D, D'	DisplayOperatorMinHeight	6
<code>\Umathfractiondel</code>	D, D'	FractionDelimiterDisplayStyleSize ⁹	delim1
<code>\Umathfractiondenomdown</code>	T, T', S, S', SS, SS'	FractionDelimiterSize ⁹	delim2
<code>\Umathfractiondenomvgap</code>	D, D'	FractionDenominatorDisplayStyleShiftDown	denom1
	T, T', S, S', SS, SS'	FractionDenominatorShiftDown	denom2
<code>\Umathfractionnumup</code>	D, D'	FractionDenominatorDisplayStyleGapMin	3*default_rule_thickness
	T, T', S, S', SS, SS'	FractionDenominatorGapMin	default_rule_thickness
<code>\Umathfractionnumvgap</code>	D, D'	FractionNumeratorDisplayStyleShiftUp	num1
	T, T', S, S', SS, SS'	FractionNumeratorShiftUp	num2
<code>\Umathfractionnumvgap</code>	D, D'	FractionNumeratorDisplayStyleGapMin	3*default_rule_thickness
	T, T', S, S', SS, SS'	FractionNumeratorGapMin	default_rule_thickness
<code>\Umathfractionrule</code>	-	FractionRuleThickness	default_rule_thickness
<code>\Umathskewedfractionhgap</code>	-	SkewedFractionHorizontalGap	math_quad/2
<code>\Umathskewedfractionvgap</code>	-	SkewedFractionVerticalGap	math_x_height



<code>\Umathlimitabovebgap</code>	-	UpperLimitBaselineRiseMin	<code>big_op_spacing3</code>
<code>\Umathlimitabovekern</code>	-	0^1	<code>big_op_spacing5</code>
<code>\Umathlimitabovevgap</code>	-	UpperLimitGapMin	<code>big_op_spacing1</code>
<code>\Umathlimitbelowbgap</code>	-	LowerLimitBaselineDropMin	<code>big_op_spacing4</code>
<code>\Umathlimitbelowkern</code>	-	0^1	<code>big_op_spacing5</code>
<code>\Umathlimitbelowvgap</code>	-	LowerLimitGapMin	<code>big_op_spacing2</code>
<code>\Umathoverdelimitervgap</code>	-	StretchStackGapBelowMin	<code>big_op_spacing1</code>
<code>\Umathoverdelimiterbgap</code>	-	StretchStackTopShiftUp	<code>big_op_spacing3</code>
<code>\Umathunderdelimitervgap</code>	-	StretchStackGapAboveMin	<code>big_op_spacing2</code>
<code>\Umathunderdelimeterbgap</code>	-	StretchStackBottomShiftDown	<code>big_op_spacing4</code>
<code>\Umathoverbarkern</code>	-	OverbarExtraAscender	<code>default_rule_thickness</code>
<code>\Umathoverbarrule</code>	-	OverbarRuleThickness	<code>default_rule_thickness</code>
<code>\Umathoverbarvgap</code>	-	OverbarVerticalGap	$3*\text{default_rule_thickness}$
<code>\Umathquad</code>	-	$\langle\text{font_size}(f)\rangle^1$	<code>math_quad</code>
<code>\Umathradicalkern</code>	-	RadicalExtraAscender	<code>default_rule_thickness</code>
<code>\Umathradicalrule</code>	-	RadicalRuleThickness	$\langle\text{not set}\rangle^2$
<code>\Umathradicalvgap</code>	D, D'	RadicalDisplayStyleVerticalGap	$(\text{default_rule_thickness} +$
	T, T', S, S', SS, SS'	RadicalVerticalGap	$(\text{abs}(\text{math_x_height})/4))^3$ $(\text{default_rule_thickness} +$
			$(\text{abs}(\text{default_rule_thickness})/4))^3$
<code>\Umathradicaldegreebefore</code>	-	RadicalKernBeforeDegree	$\langle\text{not set}\rangle^2$
<code>\Umathradicaldegreeafter</code>	-	RadicalKernAfterDegree	$\langle\text{not set}\rangle^2$
<code>\Umathradicaldegreeraise</code>	-	RadicalDegreeBottomRaisePercent	$\langle\text{not set}\rangle^{2,7}$
<code>\Umathspaceafterscript</code>	-	SpaceAfterScript	<code>script_space</code> ⁴
<code>\Umathstackdenomdown</code>	D, D'	StackBottomDisplayStyleShiftDown	<code>denom1</code>
	T, T', S, S', SS, SS'	StackBottomShiftDown	<code>denom2</code>
<code>\Umathstacknumup</code>	D, D'	StackTopDisplayStyleShiftUp	<code>num1</code>
	T, T', S, S', SS, SS'	StackTopShiftUp	<code>num3</code>
<code>\Umathstackvgap</code>	D, D'	StackDisplayStyleGapMin	$7*\text{default_rule_thickness}$
	T, T', S, S', SS, SS'	StackGapMin	$3*\text{default_rule_thickness}$
<code>\Umathsubshiftdown</code>	-	SubscriptShiftDown	<code>sub1</code>
<code>\Umathsubshiftdrop</code>	-	SubscriptBaselineDropMin	<code>sub_drop</code>
<code>\Umathsubsupshiftdown</code>	-	SubscriptShiftDownWithSuperscript ⁸ or SubscriptShiftDown	<code>sub2</code>
<code>\Umathsubtopmax</code>	-	SubscriptTopMax	$(\text{abs}(\text{math_x_height} * 4) / 5)$
<code>\Umathsubsupvgap</code>	-	SubSuperscriptGapMin	$4*\text{default_rule_thickness}$
<code>\Umathsupbottommin</code>	-	SuperscriptBottomMin	$(\text{abs}(\text{math_x_height}) / 4)$
<code>\Umathsupshiftdrop</code>	-	SuperscriptBaselineDropMax	<code>sup_drop</code>
<code>\Umathsupshiftp</code>	D	SuperscriptShiftUp	<code>sup1</code>
	T, S, SS,	SuperscriptShiftUp	<code>sup2</code>
	D', T', S', SS'	SuperscriptShiftUpCramped	<code>sup3</code>
<code>\Umathsupsubbottommax</code>	-	SuperscriptBottomMaxWithSubscript	$(\text{abs}(\text{math_x_height} * 4) / 5)$
<code>\Umathunderbarkern</code>	-	UnderbarExtraDescender	<code>default_rule_thickness</code>
<code>\Umathunderbarrule</code>	-	UnderbarRuleThickness	<code>default_rule_thickness</code>
<code>\Umathunderbarvgap</code>	-	UnderbarVerticalGap	$3*\text{default_rule_thickness}$
<code>\Umathconnectoroverlapmin</code>	-	MinConnectorOverlap	0^5



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: tfm fonts do not set `\Umathradicalrule` because \TeX 82 uses the height of the radical instead. When this parameter is indeed not set when Lua \TeX has to typeset a radical, a backward compatibility mode will kick in that assumes that an oldstyle \TeX font is used. Also, they do not set `\Umathradicaldegreebefore`, `\Umathradicaldegreearafter`, and `\Umathradicaldegreeraise`. These are then automatically initialized to $5/18$ quad, $-10/18$ quad, and 60.

Note 3: If tfm fonts are used, then the `\Umathradicalvgap` is not set until the first time Lua \TeX has to typeset a formula because this needs parameters from both family 2 and family 3. This provides a partial backward compatibility with \TeX 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 Lua \TeX has to typeset a formula. This provides some backward compatibility with \TeX 82. But once the `\Umathspaceafterscript` is set, `\scriptspace` will never be looked at again.

Note 5: Traditional tfm fonts set `\Umathconnectoroverlapmin` to zero because \TeX 82 always stacks extensibles without any overlap.

Note 6: The `\Umathoperatorsiz` 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 \TeX 82.

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

5.7 Math spacing setting

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 \TeX book). 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>\Umathopopspacing</code>
<code>\Umathordopspacing</code>	<code>\Umathopbinspacing</code>
<code>\Umathordbinspacing</code>	<code>\Umathoprelspacing</code>
<code>\Umathordrelspacing</code>	<code>\Umathopopenspacing</code>
<code>\Umathordopenspacing</code>	<code>\Umathopclosespacing</code>
<code>\Umathordclosespacing</code>	<code>\Umathoppunctspacing</code>
<code>\Umathordpunctspacing</code>	<code>\Umathopinnerspacing</code>
<code>\Umathordinnerspacing</code>	<code>\Umathbinordspacing</code>
<code>\Umathopordspacing</code>	<code>\Umathbinopspacing</code>



<code>\Umathbinbinspacing</code>	<code>\Umathcloseopspacing</code>
<code>\Umathbinrelspacing</code>	<code>\Umathclosebinspacing</code>
<code>\Umathbinopenspacing</code>	<code>\Umathcloserelspacing</code>
<code>\Umathbinclosespacing</code>	<code>\Umathcloseopenspacing</code>
<code>\Umathbinpunctspacing</code>	<code>\Umathcloseclosespacing</code>
<code>\Umathbininnerspacing</code>	<code>\Umathclosepunctspacing</code>
<code>\Umathrelordspacing</code>	<code>\Umathcloseinnerspacing</code>
<code>\Umathrelopspacing</code>	<code>\Umathpunctordspacing</code>
<code>\Umathrelbinspacing</code>	<code>\Umathpunctopspacing</code>
<code>\Umathrelrelspacing</code>	<code>\Umathpunctbinspacing</code>
<code>\Umathrelopenspacing</code>	<code>\Umathpunctrelspacing</code>
<code>\Umathrelclosespacing</code>	<code>\Umathpunctopenspacing</code>
<code>\Umathrelpunctspacing</code>	<code>\Umathpunctclosespacing</code>
<code>\Umathrelinnerspacing</code>	<code>\Umathpunctpunctspacing</code>
<code>\Umathopenordspacing</code>	<code>\Umathpunctinnerspacing</code>
<code>\Umathopenopspacing</code>	<code>\Umathinnerordspacing</code>
<code>\Umathopenbinspacing</code>	<code>\Umathinneropspacing</code>
<code>\Umathopenrelspacing</code>	<code>\Umathinnerbinspacing</code>
<code>\Umathopenopenspacing</code>	<code>\Umathinnerrelspacing</code>
<code>\Umathopenclosespacing</code>	<code>\Umathinneropenspacing</code>
<code>\Umathopenpunctspacing</code>	<code>\Umathinnerclosespacing</code>
<code>\Umathopeninnerspacing</code>	<code>\Umathinnerpunctspacing</code>
<code>\Umathcloseordspacing</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 especially. 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.

5.8 Math accent handling

Lua`TEX` 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 T_EX82.

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 T_EXbook), but if value that 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 nuclues. 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.

5.9 Math root extension

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`.

5.10 Math kerning in 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, LuaTeX 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 LuaTeX 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).

5.11 Scripts on horizontally extensible items like arrows

The primitives `\Uoverdelimit` and `\Uunderdelimit` allow the placement of a subscript or superscript on an automatically extensible item and `\Udelimiterover` and `\Udelimitunder` 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}}$
\Udelimiterover 0 "2194 {\hbox{\strut delimiterover}}$
\Udelimitunder 0 "2194 {\hbox{\strut delimiterunder}}$
```

will render this:

$$\begin{array}{ccccccc}
 \overdelimit & & & & & & \\
 \longleftarrow & & \longleftrightarrow & & \longrightarrow & & \\
 \underdelimit & \longleftarrow & & \longrightarrow & \underdelimit & \longleftarrow & \underdelimit
 \end{array}$$

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 `\Uhexextendible` is available that can be used like this:

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

This will render this:



Here you can also pass options, like:

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

This gives:



LuaTeX internally uses a structure that supports OpenType 'MathVariants' as well as tfm 'extendible recipes'.

5.12 Extracting values

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 came 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])}}
```

5.13 fractions

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 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>	$x + \frac{a}{b} + x$	$x + \frac{1}{2} + x$	$x + \left(\frac{a}{b}\right) + x$	$x + \left(\frac{1}{2}\right) + x$

5.14 Other Math changes

5.14.1 Verbose versions of single-character math commands

LuaTeX 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).

5.14.2 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.

5.15 Math surrounding skips

Inline math is surrounded by (optional) `\mathsurround` spacing but that is 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.

5.15.1 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 LuaT_EX we provide a bit more control by three variants that supporting 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.

5.15.2 Fixed scripts

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

```
d \Umathsubshiftdown
u \Umathsupshiftp
s \Umathsubsupshiftdown
```

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	
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.

5.15.3 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.

5.15.4 Math options


The logic in the math engine is rather complex and there are often no universal solutions (read: what works out well for one font, fails for another). Therefore some variations in the implementation will be driven by options for which a new primitive `\mathoption` has been introduced (so that we don't end up with many new commands). The approach of options also permits us to see what effect a specific solution has.

5.15.4.1 `\mathoption noitaliccompensation`

This option compensates placement for characters with a built-in italic correction.

```
{\showboxes\int}\quad
{\showboxes\int_{|}^{|}}\quad
{\showboxes\int\limits_{|}^{|}}
```

Gives (with computer modern that has such italics):



5.15.4.2 `\mathoption nocharitalic`

When two characters follow each other italic correction can interfere. The following example shows what this option does:

```
\catcode"1D443=11
```



```
\catcode"1D444=11
\catcode"1D445=11
P( PP PQR
```

Gives (with computer modern that has such italics):

$P(PPPQR$	$P(PPPQR$
0:inline	0:display
$P(PPPQR$	$P(PPPQR$
1:inline	1:display

5.15.4.3 `\mathoption useoldfractionsaling`

This option has been introduced as solution for tracker item 604 for fuzzy cases around either or not present fraction related settings for new fonts.



6 Nodes

6.1 LUA node representation

\TeX 's nodes are represented in Lua as userdata object with a variable set of fields. In the following syntax tables, such the type of such a userdata object is represented as `<node>`.

The current return value of `node.types()` is: `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), `unset` (15), `style` (16), `choice` (17), `noad` (18), `radical` (19), `fraction` (20), `accent` (21), `fence` (22), `math_char` (23), `sub_box` (24), `sub_mlist` (25), `math_text_char` (26), `delim` (27), `margin_kern` (28), `glyph` (29), `align_record` (30), `pseudo_file` (31), `pseudo_line` (32), `page_insert` (33), `split_insert` (34), `expr_stack` (35), `nested_list` (36), `span` (37), `attribute` (38), `glue_spec` (39), `attribute_list` (40), `temp` (41), `align_stack` (42), `movement_stack` (43), `if_stack` (44), `unhyphenated` (45), `hyphenated` (46), `delta` (47), `passive` (48), `shape` (49).

The `\lastnodetype` primitive is $\varepsilon\text{-}\TeX$ compliant. The valid range is still $[-1, 15]$ and glyph nodes (formerly known as char nodes) have number 0 while ligature nodes are mapped to 7. That way macro packages can use the same symbolic names as in traditional $\varepsilon\text{-}\TeX$. Keep in mind that these $\varepsilon\text{-}\TeX$ node numbers are different from the real internal ones and that there are more $\varepsilon\text{-}\TeX$ node types than 15.

You can ask for a list of fields with the `node.fields` (which takes an id) and for valid subtypes with `node.subtypes` (which takes a string because eventually we might support more used enumerations).

6.1.1 Auxiliary items

A few node-typed userdata objects do not occur in the 'normal' list of nodes, but can be pointed to from within that list. They are not quite the same as regular nodes, but it is easier for the library routines to treat them as if they were.

6.1.1.1 attribute_list and attribute items

The newly introduced attribute registers are non-trivial, because the value that is attached to a node is essentially a 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, but for completeness, here is the low-level interface.

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	<code><node></code>	pointer to the first attribute



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.

Valid fields in `attribute` items:

field	type	explanation
<code>next</code>	<code><node></code>	pointer to the next attribute
<code>number</code>	number	the attribute type id
<code>value</code>	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.

6.1.2 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
<code>next</code>	<code><node></code>	the next node in a list, or nil
<code>id</code>	number	the node's type (<code>id</code>) number
<code>subtype</code>	number	the node <code>subtype</code> identifier

The `subtype` is sometimes just a stub entry. 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.

6.1.2.1 hlist nodes

Valid fields: `attr`, `width`, `depth`, `height`, `dir`, `shift`, `glue_order`, `glue_sign`, `glue_set`, `head`
 Id: 0

field	type	explanation
<code>subtype</code>	number	0 = unknown origin, 1 = created by linebreaking, 2 = explicit box command, 3 = paragraph indentation box, 4 = alignment column or row, 5 = alignment cell 6 = equation 7 = equation number
<code>attr</code>	<code><node></code>	The head of the associated attribute list
<code>width</code>	number	
<code>height</code>	number	
<code>depth</code>	number	
<code>shift</code>	number	a displacement perpendicular to the character progression direction
<code>glue_order</code>	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
head	<node>	the first node of the body of this list
dir	string	the direction of this box, see 6.1.4.8

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.

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.

6.1.2.2 vlist nodes

Valid fields: As for hlist, except that 'shift' is a displacement perpendicular to the line progression direction, and 'subtype' only has subtypes 0, 4, and 5.

6.1.2.3 rule nodes

6.1.2.3.1 normal rules

Valid fields: `attr`, `width`, `depth`, `height`, `dir`, `index`

Id: 2

We have three subtypes. Subtype 0 is just a normal rule, a rectangle filled with ink. Subtype 1 is a reusable box, while subtype_2 is an image.

field	type	explanation
subtype	number	0 upto 3
attr	<node>	
width	number	the width of the rule; 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)
dir	string	the direction of this rule, see 6.1.4.8
index	number	an optional index that can be referred to (only for subtypes 1 and 2 and backend specific).

The subtypes 1 and 2 replace the `xform` and `ximage` whatsits and in node lists they behave like rules of subtype_0 when it comes to dimensions. Subtype 3 only has dimensions.

6.1.2.4 ins nodes

Valid fields: `attr`, `cost`, `depth`, `height`, `spec`, `head`

Id: 3

field	type	explanation
subtype	number	the insertion class
attr	<node>	



cost	number	the penalty associated with this insert
height	number	
depth	number	
head/list	<node>	the first node of the body of this insert
spec	<node>	a pointer to the <code>\splittopskip</code> glue spec

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 result. You can use `list` instead (often in functions you want to use local variable swith similar names and both names are equally sensible).

6.1.2.5 mark nodes

Valid fields: `attr`, `class`, `mark`

Id: 4

field	type	explanation
subtype	number	unused
attr	<node>	
class	number	the mark class
mark	table	a table representing a token list

6.1.2.6 adjust nodes

Valid fields: `attr`, `head`

Id: 5

field	type	explanation
subtype	number	0 = normal, 1 = 'pre'
attr	<node>	
head/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 result.

6.1.2.7 disc nodes

Valid fields: `attr`, `pre`, `post`, `replace`, `penalty`

Id: 7

field	type	explanation
subtype	number	indicates the source of a discretionary: 0 = the <code>\discretionary</code> command, 1 = the <code>\-</code> command, 2 = added automatically following a <code>-</code> , 3 = added by the hyphenation algorithm (simple), 4 = added by the hyphenation algorithm (hard, first item), 5 = added by the hyphenation algorithm (hard, second item)
attr	<node>	
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.

Warning: never assign a node list to the `pre`, `post` or `replace` field unless you are sure its internal link structure is correct, otherwise an error may be result. This limnitation will disappear in the future,

6.1.2.8 math nodes

Valid fields: `attr`, `surround`

Id: 11

field	type	explanation
subtype	number	0 = on, 1 = off
attr	<node>	
surround	number	width of the <code>\mathsurround</code> kern

6.1.2.9 glue nodes

Skips are about the only type of data objects in traditional T_EX that are not a simple value. The structure that represents the glue components of a skip is called a `glue_spec`, and it has the following accessible fields:

key	type	explanation
width	number	
stretch	number	
stretch_order	number	
shrink	number	
shrink_order	number	
writable	boolean	If this is true, you can’t assign to this <code>glue_spec</code> because it is one of the preallocated special cases.

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.

A spec node is normally references to from a glue node:

Valid fields: `attr`, `spec`, `leader`

Id: 12

field	type	explanation
subtype	number	0 = <code>\skip</code> , 1-18 = internal glue parameters, 100-103 = ‘leader’ subtypes



```

attr      <node>
spec      <node>  pointer to a glue_spec item
leader    <node>  pointer to a box or rule for leaders

```

The indirect spec approach is an optimization in the original T_EX code. First of all it saves quite some memory because all these spaces that become glue now share the same specification, and zero testing is also faster because only the pointer has to be checked.

The exact meanings of the subtypes are as follows:

```

1  \lineskip
2  \baselineskip
3  \parskip
4  \abovedisplayskip
5  \belowdisplayskip
6  \abovedisplayshortskip
7  \belowdisplayshortskip
8  \leftskip
9  \rightskip
10 \topskip
11 \splittopskip
12 \tabskip
13 \spaceskip
14 \xspaceskip
15 \parfillskip
16 \mathsurroundskip
17 \thinmuskip
18 \medmuskip
19 \thickmuskip
98 conditional math skip
99 mu\glue
100 \leaders
101 \cleaders
102 \xleaders
103 \gleaders

```

A regular word space also results in a `spaceskip` subtype (this used to be a `userskip` with subtype zero).

For convenience we provide access to the spec fields directly so that you can avoid the spec lookup. So, the following fields can also be queried or set. When you set a field and no spec is set, a spec will automatically be created.

key	type	explanation
width	number	
stretch	number	
stretch_order	number	
shrink	number	
shrink_order	number	



When you assign the properties to a spec using the above keys the advantage is that when needed a new spec is allocated. if you access the spec node directly you can get an error message with respect to a non-writable spec node.

By using the accessors in the glue node you are more future proof as we might decide at some point to carry all information in the glue nodes themselves. Of course we can then also decide to make the spec field kind of virtual to keep compatibility (for a while).

6.1.2.10 kern nodes

Valid fields: `attr`, `kern`, `expansion_factor`

Id: 13

field	type	explanation
subtype	number	0 = from font, 1 = from <code>\kern</code> , 2 = from <code>\accent</code> , 3 = from <code>\/</code>
attr	<node>	
kern	number	

6.1.2.11 penalty nodes

Valid fields: `attr`, `penalty`

Id: 14

field	type	explanation
subtype	number	not used
attr	<node>	
penalty	number	

6.1.2.12 glyph nodes

Valid fields: `attr`, `char`, `font`, `lang`, `left`, `right`, `uchyph`, `components`, `xoffset`, `yoffset`, `width`, `height`, `depth`, `expansion_factor`

Id: 29

field	type	explanation
subtype	number	bitfield
attr	<node>	
char	number	
font	number	
lang	number	
left	number	
right	number	
uchyph	boolean	
components	<node>	pointer to ligature components
xoffset	number	
yoffset	number	
width	number	



height	number
depth	number
expansion_factor	number

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

See section 3.1 for a detailed description of the `subtype` field.

The `expansion_factor` has been introduced as part of the separation between font- and back-end. It is the result of extensive experiments with a more efficient implementation of expansion. Early versions of LuaTeX already replaced multiple instances of fonts in the backend by scaling but contrary to pdfTeX in LuaTeX 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.

6.1.2.13 margin_kern nodes

Valid fields: `attr`, `width`, `glyph`

Id: 28

field	type	explanation
subtype	number	0 = left side, 1 = right side
attr	<node>	
width	number	
glyph	<node>	

6.1.3 Math nodes

These are the so-called 'noad's and the nodes that are specifically associated with math processing. Most of these nodes contain subnodes so that the list of possible fields is actually quite small. First, the subnodes:



6.1.3.1 Math kernel subnodes

Many object fields in math mode are either simple characters in a specific family or math lists or node lists. There are four associated subnodes that represent these cases (in the following node descriptions these are indicated by the word `<kernel>`).

The `next` and `prev` fields for these subnodes are unused.

6.1.3.1.1 `math_char` and `math_text_char` subnodes

Valid fields: `attr`, `fam`, `char`

Id: 23

field	type	explanation
<code>attr</code>	<code><node></code>	
<code>char</code>	number	
<code>fam</code>	number	

The `math_char` is the simplest subnode field, it contains the character and family for a single glyph object. 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).

6.1.3.1.2 `sub_box` and `sub_mlist` subnodes

Valid fields: `attr`, `head`

Id: 24

field	type	explanation
<code>attr</code>	<code><node></code>	
<code>head</code>	<code><node></code>	

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

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 result.

6.1.3.2 Math delimiter subnode

There is a fifth subnode type that is used exclusively for delimiter fields. As before, the `next` and `prev` fields are unused.

6.1.3.2.1 `delim` subnodes

Valid fields: `attr`, `small_fam`, `small_char`, `large_fam`, `large_char`

Id: 27



field	type	explanation
attr	<node>	
small_char	number	
small_fam	number	
large_char	number	
large_fam	number	

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`.

6.1.3.3 Math core nodes

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`, `vcent`. These all have the same fields, and they are combined into a single node type with separate subtypes for differentiation.

6.1.3.3.1 simple nodes

Valid fields: `attr`, `nucleus`, `sub`, `sup`

Id: 18

field	type	explanation
subtype	number	see below
attr	<node>	
nucleus	<kernel>	
sub	<kernel>	
sup	<kernel>	

Operators are a bit special because they occupy three subtypes. `subtype`.

number	node subtype
0	Ord
1	Op: <code>\displaylimits</code>
2	Op: <code>\limits</code>
3	Op: <code>\nolimits</code>
4	Bin
5	Rel
6	Open
7	Close
8	Punct
9	Inner
10	Under
11	Over
12	Vcent



6.1.3.3.2 accent nodes

Valid fields: `attr`, `nucleus`, `sub`, `sup`, `accent`, `bot_accent`, `top_accent`, `overlay_accent`

Id: 21

field	type	explanation
subtype	number	the first bit is used for a fixed top accent flag (if the <code>accent</code> field is present), the second bit for a fixed bottom accent flag (if the <code>bot_accent</code> field is present); example: the actual value 3 means: do not stretch either accent
attr	<node>	
nucleus	<kernel>	
sub	<kernel>	
sup	<kernel>	
accent	<kernel>	
bot_accent	<kernel>	

6.1.3.3.3 style nodes

Valid fields: `attr`, `style`

Id: 16

field	type	explanation
style	string	contains the style

There are eight possibilities for the string value: one of 'display', 'text', 'script', or 'scriptscript'. Each of these can have a trailing ' to signify 'cramped' styles.

6.1.3.3.4 choice nodes

Valid fields: `attr`, `display`, `text`, `script`, `scriptscript`

Id: 17

field	type	explanation
attr	<node>	
display	<node>	
text	<node>	
script	<node>	
scriptscript	<node>	

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

6.1.3.3.5 radical nodes

Valid fields: `attr`, `nucleus`, `sub`, `sup`, `left`, `degree`

Id: 19



field	type	explanation
attr	<node>	
nucleus	<kernel>	
sub	<kernel>	
sup	<kernel>	
left	<delim>	
degree	<kernel>	Only set by <code>\Uroot</code>

A 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 may be result.

The radical noad is also used for under- and overdelimiters, which is indicated by the subtypes:

- 0 `\radical`
- 1 `\Uradical`
- 2 `\Uroot`
- 3 `\Uunderdelimiter`
- 4 `\Uoverdelimiter`
- 5 `\Udelimiterunder`
- 6 `\Udelimiterover`

6.1.3.3.6 fraction nodes

Valid fields: attr, width, num, denom, left, right

Id: 20

field	type	explanation
attr	<node>	
width	number	
num	<kernel>	
denom	<kernel>	
left	<delim>	
right	<delim>	

A 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 may be result.

6.1.3.3.7 fence nodes

Valid fields: attr, delim

Id: 22

field	type	explanation
subtype	number	1 = <code>\left</code> , 2 = <code>\middle</code> , 3 = <code>\right</code>
attr	<node>	
delim	<delim>	



6.1.4 whatsit nodes

Whatsit nodes come in many subtypes that you can ask for by running `node.whatsits()`: `open` (0), `write` (1), `close` (2), `special` (3), `save_pos` (6), `late_lua` (7), `user_defined` (8), `pdf_literal` (16), `pdf_refobj` (17), `pdf_annot` (18), `pdf_start_link` (19), `pdf_end_link` (20), `pdf_dest` (21), `pdf_action` (22), `pdf_thread` (23), `pdf_start_thread` (24), `pdf_end_thread` (25), `pdf_thread_data` (26), `pdf_link_data` (27), `pdf_colorstack` (28), `pdf_setmatrix` (29), `pdf_save` (30), `pdf_restore` (31), `fake` (100).

6.1.4.1 open nodes

Valid fields: `attr`, `stream`, `name`, `area`, `ext`
Id: 8, 0

field	type	explanation
<code>attr</code>	<node>	
<code>stream</code>	number	T _E X's stream id number
<code>name</code>	string	file name
<code>ext</code>	string	file extension
<code>area</code>	string	file area (this may become obsolete)

6.1.4.2 write nodes

Valid fields: `attr`, `stream`, `data`
Id: 8, 1

field	type	explanation
<code>attr</code>	<node>	
<code>stream</code>	number	T _E X's stream id number
<code>data</code>	table	a table representing the token list to be written

6.1.4.3 close nodes

Valid fields: `attr`, `stream`
Id: 8, 2

field	type	explanation
<code>attr</code>	<node>	
<code>stream</code>	number	T _E X's stream id number

6.1.4.4 special nodes

Valid fields: `attr`, `data`
Id: 8, 3

field	type	explanation
<code>attr</code>	<node>	
<code>data</code>	string	the <code>\special</code> information



6.1.4.5 boundary nodes

Valid fields: `attr`, `value`

Id: 6

This node relates to the `\noboundary` primitive but you can use it for your own purpose too, in which case `\boundary` can come in handy.

6.1.4.6 language nodes

LuaTeX does not have language whatsits any more. All language information is already present inside the glyph nodes themselves. This whatsit subtype will be removed in the next release.

6.1.4.7 local_par nodes

Valid fields: `attr`, `pen_inter`, `pen_broken`, `dir`, `box_left`, `box_left_width`, `box_right`, `box_right_width`

Id: 9

field	type	explanation
<code>attr</code>	<code><node></code>	
<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 6.1.4.8
<code>box_left</code>	<code><node></code>	the <code>\localleftbox</code>
<code>box_left_width</code>	number	width of the <code>\localleftbox</code>
<code>box_right</code>	<code><node></code>	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 be result.

6.1.4.8 dir nodes

Valid fields: `attr`, `dir`, `level`, `dvi_ptr`, `dvi_h`

Id: 10

field	type	explanation
<code>attr</code>	<code><node></code>	
<code>dir</code>	string	the direction (but see below)
<code>level</code>	number	nesting level of this direction whatsit
<code>dvi_ptr</code>	number	a saved dvi buffer byte offset
<code>dir_h</code>	number	a saved dvi position

A note on `dir` strings. Direction specifiers are three-letter combinations of T, B, R, and L.

These are built up out of three separate items:



- the first is the direction of the ‘top’ of paragraphs.
- the second is the direction of the ‘start’ of lines.
- the third is the direction of the ‘top’ of glyphs.

However, only four combinations are accepted: `TLT`, `TRT`, `RTT`, and `LTL`.

Inside actual `dir` whatsit nodes, the representation of `dir` is not a three-letter but a four-letter combination. The first character in this case is always either `+` or `-`, indicating whether the value is pushed or popped from the direction stack.

6.1.4.9 pdf_literal nodes

Valid fields: `attr`, `mode`, `data`

Id: 8, 16

field	type	explanation
<code>attr</code>	<code><node></code>	
<code>mode</code>	number	the ‘mode’ setting of this literal
<code>data</code>	string	the <code>\pdfliteral</code> information

Mode values:

value	corresponding <code>\pdfTeX</code> keyword
0	<code>setorigin</code>
1	<code>page</code>
2	<code>direct</code>

6.1.4.10 pdf_refobj nodes

Valid fields: `attr`, `objnum`

Id: 8, 17

field	type	explanation
<code>attr</code>	<code><node></code>	
<code>objnum</code>	number	the referenced pdf object number

6.1.4.11 pdf_annot nodes

Valid fields: `attr`, `width`, `depth`, `height`, `objnum`, `data`

Id: 8, 18

field	type	explanation
<code>attr</code>	<code><node></code>	
<code>width</code>	number	
<code>height</code>	number	
<code>depth</code>	number	
<code>objnum</code>	number	the referenced pdf object number
<code>data</code>	string	the annotation data



6.1.4.12 pdf_start_link nodes

Valid fields: attr, width, depth, height, objnum, link_attr, action

Id: 8, 19

field	type	explanation
attr	<node>	
width	number	
height	number	
depth	number	
objnum	number	the referenced pdf object number
link_attr	table	the link attribute token list
action	<node>	the action to perform

6.1.4.13 pdf_end_link nodes

Valid fields: attr

Id: 8, 20

field	type	explanation
attr	<node>	

6.1.4.14 pdf_dest nodes

Valid fields: attr, width, depth, height, named_id, dest_id, dest_type, xyz_zoom, objnum

Id: 8, 21

field	type	explanation
attr	<node>	
width	number	
height	number	
depth	number	
named_id	number	is the dest_id a string value?
dest_id	number	the destination id
	string	the destination name
dest_type	number	type of destination
xyz_zoom	number	
objnum	number	the pdf object number

6.1.4.15 pdf_action nodes

Valid fields: action_type, named_id, action_id, file, new_window, data, ref_count

Id: 8, 22

These are a special kind of item that only appears inside pdf start link objects.

field	type	explanation
action_type	number	



action_id	number or string	
named_id	number	
file	string	
new_window	number	
data	string	
ref_count	number	read-only

6.1.4.16 pdf_thread nodes

Valid fields: attr, width, depth, height, named_id, thread_id, thread_attr
 Id: 8, 23

field	type	explanation
attr	<node>	
width	number	
height	number	
depth	number	
named_id	number	is tread_id a string value?
tread_id	number	the thread id
	string	the thread name
thread_attr	number	extra thread information

6.1.4.17 pdf_start_thread nodes

Valid fields: attr, width, depth, height, named_id, thread_id, thread_attr
 Id: 8, 24

field	type	explanation
attr	<node>	
width	number	
height	number	
depth	number	
named_id	number	is tread_id a string value?
tread_id	number	the thread id
	string	the thread name
thread_attr	number	extra thread information

6.1.4.18 pdf_end_thread nodes

Valid fields: attr
 Id: 8, 25

field	type	explanation
attr	<node>	

6.1.4.19 save_pos nodes

Valid fields: attr
 Id: 8, 6



field	type	explanation
attr	<node>	

6.1.4.20 late_lua nodes

Valid fields: attr, reg, data, name, string
 Id: 8, 7

field	type	explanation
attr	<node>	
data	string	data to execute
string	string	data to execute
name	string	the name to use for Lua error reporting

The difference between `data` and `string` is that on assignment, the `data` field is converted to a token list, cf. use as `\latelua`. The `string` version is treated as a literal string.

6.1.4.21 pdf_colorstack nodes

Valid fields: attr, stack, cmd, data
 Id: 8, 28

field	type	explanation
attr	<node>	
stack	number	colorstack id number
command	number	command to execute
data	string	data

6.1.4.22 pdf_setmatrix nodes

Valid fields: attr, data
 Id: 8, 29

field	type	explanation
attr	<node>	
data	string	data

6.1.4.23 pdf_save nodes

Valid fields: attr
 Id: 8, 30

field	type	explanation
attr	<node>	

6.1.4.24 pdf_restore nodes

Valid fields: attr
 Id: 8, 31



field	type	explanation
attr	<node>	

6.1.4.25 user_defined nodes

User-defined whatsit nodes can only be created and handled from Lua code. In effect, they are an extension to the extension mechanism. The LuaTeX engine will simply step over such whatsites without ever looking at the contents.

Valid fields: attr, user_id, type, value

Id: 8, 8

field	type	explanation
attr	<node>	
user_id	number	id number
type	number	type of the value
value	number	
	string	
	<node>	
	table	

The type can have one of five distinct values:

value	explanation
97	the value is an attribute node list
100	the value is a number
110	the value is a node list
115	the value is a string
116	the value is a token list in Lua table form

6.2 Two access models

Deep down in TeX a node has a number which is a numeric entry in a memory table. In fact, this model, where TeX manages memory is real fast and one of the reasons why plugging in callbacks that operate on nodes is quite fast. So, 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 our advice is: use the indexed (table) approach when possible and investigate the direct one when speed might be an issue. For that reason we 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 we're accessing nodes 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 LuaTeX 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
```

If performance matters you can use an function instead:

<code>getnext</code>	parsing nodelist always involves this one
<code>getprev</code>	used less but is logical companion to <code>getnext</code>
<code>getboth</code>	returns the next and prev pointer of a node
<code>getid</code>	consulted a lot
<code>getsubtype</code>	consulted less but also a topper
<code>getfont</code>	used a lot in OpenType handling (glyph nodes are consulted a lot)
<code>getchar</code>	idem and also in other places
<code>getdisc</code>	returns the <code>pre</code> , <code>post</code> and <code>replace</code> fields and optionally when true is passed also the tail fields.
<code>getList</code>	we often parse nested lists so this is a convenient one too (only works for <code>hlist</code> and <code>vlist</code> !)
<code>getleader</code>	comparable to <code>list</code> , seldom used in TeX (but needs frequent consulting like lists; leaders could have been made a dedicated node type)
<code>getField</code>	generic getter, sufficient for the rest (other field names are often shared so a specific getter makes no sense then)

The direct variants also have setters, where the discretionary setter takes three (optional) arguments plus an optional fourth indicating the subtype.

It doesn't make sense to add more. Profiling demonstrated that these fields can get accesses way more times than other fields. Even in complex documents, many node and fields types never get seen, or seen only a few times. Most functions in the `node` namespace have a companion in `node.direct`, but of course not the ones that don't deal with nodes themselves. The following table summarized this:

function	node	direct
-----------------	-------------	---------------



copy_list	+	+
copy	+	+
count	+	+
current_attr	+	+
dimensions	+	+
do_ligature_n	+	+
effective_glue	+	+
end_of_math	+	+
family_font	+	-
fields	+	-
first_character	+	-
first_glyph	+	+
flush_list	+	+
flush_node	+	+
free	+	+
getboth	+	+
getbox	-	+
getchar	+	+
getdisc	+	+
getfield	+	+
getfont	+	+
getid	+	+
getleader	+	+
getlist	+	+
getnext	+	+
getprev	+	+
getsubtype	+	+
has_attribute	+	+
has_field	+	+
has_glyph	+	+
hpack	+	+
id	+	-
insert_after	+	+
insert_before	+	+
is_char	+	+
is_glyph	+	+
is_direct	-	+
is_node	+	+
kerning	+	+
last_node	+	+
length	+	+
ligaturing	+	+
mlist_to_hlist	+	-
new	+	+
next	+	-



<code>prev</code>	+	-
<code>protect_glyph</code>	+	+
<code>protect_glyphs</code>	+	+
<code>protrusion_skippable</code>	+	+
<code>remove</code>	+	+
<code>set_attribute</code>	+	+
<code>setboth</code>	+	+
<code>setbox</code>	+	+
<code>setchar</code>	+	+
<code>setdisc</code>	+	+
<code>setfield</code>	+	+
<code>setlink</code>	+	+
<code>setnext</code>	+	+
<code>setprev</code>	+	+
<code>slide</code>	+	+
<code>subtype</code>	+	-
<code>subtypes</code>	+	-
<code>tail</code>	+	+
<code>todirect</code>	+	+
<code>tonode</code>	+	+
<code>tostring</code>	+	+
<code>traverse_id</code>	+	+
<code>traverse_char</code>	+	+
<code>traverse</code>	+	+
<code>types</code>	+	-
<code>type</code>	+	-
<code>unprotect_glyphs</code>	+	+
<code>unset_attribute</code>	+	+
<code>usedlist</code>	+	+
<code>vpack</code>	+	+
<code>whatsits</code>	+	-
<code>whatsitsubtypes</code>	+	-
<code>write</code>	+	+

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 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.



7 L^AT_EX LUA Libraries

The implied use of the built-in Lua modules `epdf`, `fontloader`, `mplib`, and `pdfscanner` is deprecated. If you want to use these, please start your source file with a proper `require` line. In the future, Lua_TE_X will switch to loading these modules on demand.

The interfacing between T_EX and Lua is facilitated by a set of library modules. The Lua libraries in this chapter are all defined and initialized by the Lua_TE_X executable. Together, they allow Lua scripts to query and change a number of T_EX's internal variables, run various internal T_EX functions, and set up Lua_TE_X's hooks to execute Lua code.

The following sections are in alphabetical order. For any callback (and manipulation of nodes) the following is true: you have a lot of freedom which also means that you can mess up the node lists and nodes themselves. So, a bit of defensive programming doesn't hurt. A crash can happen when you spoil things or when Lua_TE_X can recognize the issue, a panic exit will happen. Don't bother the team with such issues.

7.1 The callback library

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 mosty 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_TE_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_TE_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, `callback.find` returns `nil`.

7.1.1 File discovery callbacks

The behaviour documented in this subsection is considered stable in the sense that there will not be backward-incompatible changes any more.

7.1.1.1 find_read_file and find_write_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 T_EX'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 T_EX, 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 LuaT_EX 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.

7.1.1.2 find_font_file

Your callback function should have the following conventions:

```
<string> actual_name = function (<string> asked_name)
```

The `asked_name` is an otf or tfm font metrics file.

Return `nil` if the file cannot be found.

7.1.1.3 find_output_file

Your callback function should have the following conventions:



```
<string> actual_name = function (<string> asked_name)
```

The `asked_name` is the pdf or dvi file for writing.

7.1.1.4 `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).

7.1.1.5 `find_vf_file`

Like `find_font_file`, but for virtual fonts. This applies to both Aleph's ovf files and traditional Knuthian vf files.

7.1.1.6 `find_map_file`

Like `find_font_file`, but for map files.

7.1.1.7 `find_enc_file`

Like `find_font_file`, but for enc files.

7.1.1.8 `find_sfd_file`

Like `find_font_file`, but for subfont definition files.

7.1.1.9 `find_pk_file`

Like `find_font_file`, but for pk bitmap files. This callback takes two arguments: `name` and `dpi`. In your callback you can decide to look for:

```
<base res>dpi/<fontname>.<actual res>pk
```

but other strategies are possible. It is up to you to find a 'reasonable' bitmap file to go with that specification.

7.1.1.10 `find_data_file`

Like `find_font_file`, but for embedded files (`\pdfobj file '...'`).

7.1.1.11 `find_opentype_file`

Like `find_font_file`, but for OpenType font files.



7.1.1.12 `find_truetype_file` and `find_type1_file`

Your callback function should have the following conventions:

```
<string> actual_name = function (<string> asked_name)
```

The `asked_name` is a font file. This callback is called while LuaTeX is building its internal list of needed font files, so the actual timing may surprise you. Your return value is later fed back into the matching `read_file` callback.

Strangely enough, `find_type1_file` is also used for OpenType (otf) fonts.

7.1.1.13 `find_image_file`

Your callback function should have the following conventions:

```
<string> actual_name = function (<string> asked_name)
```

The `asked_name` is an image file. Your return value is used to open a file from the hard disk, so make sure you return something that is considered the name of a valid file by your operating system.

7.1.2 File reading callbacks

The behavior documented in this subsection is considered stable in the sense that there will not be backward-incompatible changes any more.

7.1.2.1 `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.



7.1.2.1.1 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.

7.1.2.1.2 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.

7.1.2.2 General file readers

There is a set of callbacks for the loading of binary data files. These all use the same interface:

```
function(<string> name)
    return <boolean> success, <string> data, <number> data_size
end
```

The `name` will normally be a full path name as it is returned by either one of the file discovery callbacks or the internal version of `kpse.find_file()`.

`success`

Return `false` when a fatal error occurred (e.g. when the file cannot be found, after all).

`data`

The bytes comprising the file.

`data_size`

The length of the `data`, in bytes.

Return an empty string and zero if the file was found but there was a reading problem.

The list of functions is as follows:

<code>read_font_file</code>	ofm or tfm files
<code>read_vf_file</code>	virtual fonts
<code>read_map_file</code>	map files
<code>read_enc_file</code>	encoding files
<code>read_sfd_file</code>	subfont definition files
<code>read_pk_file</code>	pk bitmap files
<code>read_data_file</code>	embedded files (<code>\pdfobj file ...</code>)



`read_truetype_file` TrueType font files
`read_type1_file` Type1 font files
`read_opentype_file` OpenType font files

7.1.3 Data processing callbacks

7.1.3.1 `process_input_buffer`

This callback allows you to change the contents of the line input buffer just before LuaTeX actually starts looking at it.

```
function(<string> buffer)
    return <string> adjusted_buffer
end
```

If you return `nil`, LuaTeX will pretend like your callback never happened. You can gain a small amount of processing time from that. This callback does not replace any internal code.

7.1.3.2 `process_output_buffer`

This callback allows you to change the contents of the line output buffer just before LuaTeX actually starts writing it to a file as the result of a `\write` command. It is only called for output to an actual file (that is, excluding the log, the terminal, and `\write18` calls).

```
function(<string> buffer)
    return <string> adjusted_buffer
end
```

If you return `nil`, LuaTeX will pretend like your callback never happened. You can gain a small amount of processing time from that. This callback does not replace any internal code.

7.1.3.3 `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`, LuaTeX will pretend your callback never happened. This callback does not replace any internal code.



7.1.4 Node list processing callbacks

The description of nodes and node lists is in chapter 6.

7.1.4.1 `contribute_filter`

This callback is called when LuaTeX 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.

group codes	pointer
<code>pre_box</code>	<code>contrib_head</code>
<code>pre_adjust_tail</code>	<code>pre_adjust_head</code>
<code>just</code>	<code>just_box</code>
<code>adjust_tail</code>	<code>adjust_head</code>

7.1.4.2 `buildpage_filter`

This callback is called whenever LuaTeX 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 TeX's state is with respect to the 'current page'. The possible values are:

value	explanation
<code>alignment</code>	a (partial) alignment is being added
<code>after_output</code>	an output routine has just finished
<code>box</code>	a typeset box is being added
<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
<code>after_display</code>	a display is finished
<code>end</code>	LuaTeX is terminating (it's all over)

This callback does not replace any internal code.

7.1.4.3 `pre_linebreak_filter`

This callback is called just before LuaTeX 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 TeX’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
<code><empty></code>	main vertical list
<code>hbox</code>	<code>\hbox</code> in horizontal mode
<code>adjusted_hbox</code>	<code>\hbox</code> in vertical mode
<code>vbox</code>	<code>\vbox</code>
<code>vtop</code>	<code>\vtop</code>
<code>align</code>	<code>\halign</code> or <code>\valign</code>
<code>disc</code>	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.

7.1.4.4 `linebreak_filter`

This callback replaces LuaTeX’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>`, LuaTeX 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’.

7.1.4.5 `append_to_vlist_filter`

This callback is called whenever LuaTeX 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`.

7.1.4.6 `post_linebreak_filter`

This callback is called just after LuaTeX 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.

7.1.4.7 `hpack_filter`

This callback is called when TeX 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.

7.1.4.8 `vpack_filter`

This callback is called when TeX 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.

7.1.4.9 `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 `overfull`, `underfull`, `loose` or `tight`. The detail is either the amount of overflow in case of `overfull`, 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 overfull rule indicator. That node will be appended to the list (just like T_EX's own rule would).

7.1.4.10 `vpack_quality`

This callback can be used to intercept the overfull 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 `overfull`, `underfull`, `loose` or `tight`. The detail is either the amount of overflow in case of `overfull`, or the badness otherwise. The head is the list that is constructed.

7.1.4.11 `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.

7.1.4.12 `pre_output_filter`

This callback is called when T_EX 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.

7.1.4.13 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.

7.1.4.14 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 LuaTeX into panic mode.

7.1.4.15 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 LuaTeX into panic mode.



7.1.4.16 `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 \TeX can get confused.

7.1.4.17 `mlist_to_hlist`

This callback replaces Lua \TeX '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.

7.1.5 Information reporting callbacks

7.1.5.1 `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.

7.1.5.2 `start_run`

```
function()
end
```

This callback replaces the code that prints Lua \TeX '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.

7.1.5.3 `stop_run`

```
function()
end
```



This callback replaces the code that prints LuaTeX's statistics and 'output written to' messages.

7.1.5.4 `start_page_number`

```
function()  
end
```

Replaces the code that prints the [and the page number at the begin of `\shipout`. This callback will also override the printing of box information that normally takes place when `\tracingoutput` is positive.

7.1.5.5 `stop_page_number`

```
function()  
end
```

Replaces the code that prints the] at the end of `\shipout`.

7.1.5.6 `show_error_hook`

```
function()  
end
```

This callback is run from inside the TeX error function, and the idea is to allow you to do some extra reporting on top of what TeX 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.

7.1.5.7 `show_error_message`

```
function()  
end
```

This callback replaces the code that prints the error message. The usual interaction after the message is not affected.

7.1.5.8 `show_lua_error_hook`

```
function()  
end
```

This callback replaces the code that prints the extra Lua error message.

7.1.5.9 `start_file`

```
function(category, filename)  
end
```



This callback replaces the code that prints Lua \TeX 's when a file is opened like (`filename` for regular files. The category is a number:

- 1 a normal data file, like a \TeX 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

7.1.5.10 `stop_file`

```
function(category)
end
```

This callback replaces the code that prints Lua \TeX 's when a file is closed like the `)` for regular files.

7.1.6 PDF-related callbacks

7.1.6.1 `finish_pdffile`

```
function()
end
```

This callback is called when all document pages are already written to the pdf file and Lua \TeX is about to finalize the output document structure. Its intended use is final update of pdf dictionaries such as `/Catalog` or `/Info`. The callback does not replace any code. There are neither arguments nor return values.

7.1.6.2 `finish_pdfpage`

```
function(shippingout)
end
```

This callback is called after the pdf page stream has been assembled and before the page object gets finalized.

7.1.7 Font-related callbacks

7.1.7.1 `define_font`

```
function(<string> name, <number> size, <number> id)
    return <table> font | <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 designsize of the font.

The `id` is the internal number assigned to the font.

The internal structure of the `font` table that is to be returned is explained in chapter 4. 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.

7.2 The `epdf` library

The `epdf` library provides Lua bindings to many pdf access functions that are defined by the poppler pdf viewer library (written in C++ by Kristian Høgsberg, based on `xpdf` by Derek Noonburg). Within LuaTeX (and pdfTeX), `xpdf` functionality is being used since long time to embed pdf files. The `epdf` library shall allow to scrutinize an external pdf file. It gives access to its document structure: catalog, cross-reference table, individual pages, objects, annotations, info, and metadata. The LuaTeX team is evaluating the possibility of reducing the binding to a basic low level pdf primitives and delegate the complete set of functions to an external shared object module.

The `epdf` library is still in alpha state: pdf access is currently read-only. It’s not yet possible to alter a pdf file or to assemble it from scratch, and many function bindings are still missing, and it is unlikely that we to support that at all. At some point we might also decide to limit the interface to a reasonable subset.

For a start, a pdf file is opened by `epdf.open()` with file name, e.g.:

```
doc = epdf.open("foo.pdf")
```

This normally returns a `PDFDoc` userdata variable; but if the file could not be opened successfully, instead of a fatal error just the value `nil` is returned.

All Lua functions in the `epdf` library are named after the poppler functions listed in the poppler header files for the various classes, e.g., files `PDFDoc.h`, `Dict.h`, and `Array.h`. These files can be found in the poppler subdirectory within the LuaTeX sources. Which functions are already implemented in the `epdf` library can be found in the LuaTeX source file `lepdfplib.cc`. For using the `epdf` library, knowledge of the pdf file architecture is indispensable.

There are many different userdata types defined by the `epdf` library, currently these are `AnnotBorderStyle`, `AnnotBorder`, `Annots`, `Annot`, `Array`, `Attribute`, `Catalog`, `Dict`, `EmbFile`, `GString`, `LinkDest`, `Links`, `Link`, `ObjectStream`, `Object`, `PDFDoc`, `PDFRectangle`, `Page`, `Ref`, `Stream`, `StructElement`, `StructTreeRoot` `TextSpan`, `XRefEntry` and `XRef`.

All these userdata names and the Lua access functions closely resemble the classes naming from the poppler header files, including the choice of mixed upper and lower case letters. The



Lua function calls use object-oriented syntax, e.g., the following calls return the `Page` object for page 1:

```
pageref = doc:getCatalog():getPageRef(1)
pageobj  = doc:getXRef():fetch(pageref.num, pageref.gen)
```

But writing such chained calls is risky, as an intermediate function may return `nil` on error; therefore between function calls there should be Lua type checks (e.g., against `nil`) done. If a non-object item is requested (e.g., a `Dict` item by calling `page:getPieceInfo()`, cf. `Page.h`) but not available, the Lua functions return `nil` (without error). If a function should return an `Object`, but it's not existing, a `Null` object is returned instead (also without error; this is in-line with poppler behavior).

All library objects have a `__gc` metamethod for garbage collection. The `__tostring` metamethod gives the type name for each object.

All object constructors:

```
<PDFDoc>      = epdf.open(<string> PDF filename)
<Annot>       = epdf.Annot(<XRef>, <Dict>, <Catalog>, <Ref>)
<Annots>      = epdf.Annots(<XRef>, <Catalog>, <Object>)
<Array>       = epdf.Array(<XRef>)
<Attribute>   = epdf.Attribute(<Type>,<Object>)| epdf.Attribute(<string>, <int>,
<Object>)
<Dict>        = epdf.Dict(<XRef>)
<Object>      = epdf.Object()
<PDFRectangle> = epdf.PDFRectangle()
```

The functions `StructElement_Type`, `Attribute_Type` and `AttributeOwner_Type` return a hash table `{<string>,<integer>}`.

Annot methods:

```
<boolean>    = <Annot>:isOk()
<Object>     = <Annot>:getAppearance()
<AnnotBorder> = <Annot>:getBorder()
<boolean>    = <Annot>:match(<Ref>)
```

AnnotBorderStyle methods:

```
<number> = <AnnotBorderStyle>:getWidth()
```

Annots methods:

```
<integer> = <Annots>:getNumAnnots()
<Annot>   = <Annots>:getAnnot(<integer>)
```

Array methods:

```
<Array>:incRef()
<Array>:decRef()
```



```

<integer> = <Array>:getLength()
           <Array>:add(<Object>)
<Object>  = <Array>:get(<integer>)
<Object>  = <Array>:getNF(<integer>)
<string>  = <Array>:getString(<integer>)

```

Attribute methods:

```

<boolean> = <Attribute>:isOk()
<integer> = <Attribute>:getType()
<integer> = <Attribute>:getOwner()
<string>  = <Attribute>:getTypeName()
<string>  = <Attribute>:getOwnerName()
<Object>  = <Attribute>:getValue()
<Object>  = <Attribute>:getDefaultValue
<string>  = <Attribute>:getName()
<integer> = <Attribute>:getRevision()
           <Attribute>:setRevision(<unsigned integer>)
<boolean> = <Attribute>:isHidden()
           <Attribute>:setHidden(<boolean>)
<string>  = <Attribute>:getFormattedValue()
<string>  = <Attribute>:setFormattedValue(<string>)

```

Catalog methods:

```

<boolean> = <Catalog>:isOk()
<integer> = <Catalog>:getNumPages()
<Page>    = <Catalog>:getPage(<integer>)
<Ref>     = <Catalog>:getPageRef(<integer>)
<string>  = <Catalog>:getBaseURI()
<string>  = <Catalog>:readMetadata()
<Object>  = <Catalog>:getStructTreeRoot()
<integer> = <Catalog>:findPage(<integer> object number, <integer> object gener-
ation)
<LinkDest> = <Catalog>:findDest(<string> name)
<Object>   = <Catalog>:getDests()
<integer>  = <Catalog>:numEmbeddedFiles()
<EmbFile> = <Catalog>:embeddedFile(<integer>)
<integer>  = <Catalog>:numJS()
<string>   = <Catalog>:getJS(<integer>)
<Object>   = <Catalog>:getOutline()
<Object>   = <Catalog>:getAcroForm()

```

EmbFile methods:

```

<string>   = <EmbFile>:name()
<string>   = <EmbFile>:description()
<integer>  = <EmbFile>:size()

```



```

<string> = <EmbFile>:modDate()
<string> = <EmbFile>:createDate()
<string> = <EmbFile>:checksum()
<string> = <EmbFile>:mimeType()
<Object> = <EmbFile>:streamObject()
<boolean> = <EmbFile>:isOk()

```

Dict methods:

```

    <Dict>:incRef()
    <Dict>:decRef()
<integer> = <Dict>:getLength()
    <Dict>:add(<string>, <Object>)
    <Dict>:set(<string>, <Object>)
    <Dict>:remove(<string>)
<boolean> = <Dict>:is(<string>)
<Object> = <Dict>:lookup(<string>)
<Object> = <Dict>:lookupNF(<string>)
<integer> = <Dict>:lookupInt(<string>, <string>)
<string> = <Dict>:getKey(<integer>)
<Object> = <Dict>:getVal(<integer>)
<Object> = <Dict>:getValNF(<integer>)
<boolean> = <Dict>:hasKey(<string>)

```

Link methods:

```

<boolean> = <Link>:isOk()
<boolean> = <Link>:inRect(<number>, <number>)

```

LinkDest methods:

```

<boolean> = <LinkDest>:isOk()
<integer> = <LinkDest>:getKind()
<string> = <LinkDest>:getKindName()
<boolean> = <LinkDest>:isPageRef()
<integer> = <LinkDest>:getPageNum()
<Ref> = <LinkDest>:getPageRef()
<number> = <LinkDest>:getLeft()
<number> = <LinkDest>:getBottom()
<number> = <LinkDest>:getRight()
<number> = <LinkDest>:getTop()
<number> = <LinkDest>:getZoom()
<boolean> = <LinkDest>:getChangeLeft()
<boolean> = <LinkDest>:getChangeTop()
<boolean> = <LinkDest>:getChangeZoom()

```

Links methods:

```

<integer> = <Links>:getNumLinks()

```



<Link> = <Links>:getLink(<integer>)

Object methods:

<Object>:initBool(<boolean>)
<Object>:initInt(<integer>)
<Object>:initReal(<number>)
<Object>:initString(<string>)
<Object>:initName(<string>)
<Object>:initNull()
<Object>:initArray(<XRef>)
<Object>:initDict(<XRef>)
<Object>:initStream(<Stream>)
<Object>:initRef(<integer> object number, <integer> object genera-

tion)

<Object>:initCmd(<string>)
<Object>:initError()
<Object>:initEOF()

<Object> = <Object>:fetch(<XRef>)

<integer> = <Object>:getType()

<string> = <Object>:getTypeName()

<boolean> = <Object>:isBool()

<boolean> = <Object>:isInt()

<boolean> = <Object>:isReal()

<boolean> = <Object>:isNum()

<boolean> = <Object>:isString()

<boolean> = <Object>:isName()

<boolean> = <Object>:isNull()

<boolean> = <Object>:isArray()

<boolean> = <Object>:isDict()

<boolean> = <Object>:isStream()

<boolean> = <Object>:isRef()

<boolean> = <Object>:isCmd()

<boolean> = <Object>:isError()

<boolean> = <Object>:isEOF()

<boolean> = <Object>:isNone()

<boolean> = <Object>:getBool()

<integer> = <Object>:getInt()

<number> = <Object>:getReal()

<number> = <Object>:getNum()

<string> = <Object>:getString()

<string> = <Object>:getName()

<Array> = <Object>:getArray()

<Dict> = <Object>:getDict()

<Stream> = <Object>:getStream()

<Ref> = <Object>:getRef()

<integer> = <Object>:getRefNum()



```

<integer> = <Object>:getRefGen()
<string>  = <Object>:getCmd()
<integer> = <Object>:arrayGetLength()
           = <Object>:arrayAdd(<Object>)
<Object>  = <Object>:arrayGet(<integer>)
<Object>  = <Object>:arrayGetNF(<integer>)
<integer> = <Object>:dictGetLength(<integer>)
           = <Object>:dictAdd(<string>, <Object>)
           = <Object>:dictSet(<string>, <Object>)
<Object>  = <Object>:dictLookup(<string>)
<Object>  = <Object>:dictLookupNF(<string>)
<string>  = <Object>:dictgetKey(<integer>)
<Object>  = <Object>:dictgetVal(<integer>)
<Object>  = <Object>:dictgetValNF(<integer>)
<boolean> = <Object>:streamIs(<string>)
           = <Object>:streamReset()
<integer> = <Object>:streamGetChar()
<integer> = <Object>:streamLookChar()
<integer> = <Object>:streamGetPos()
           = <Object>:streamSetPos(<integer>)
<Dict>    = <Object>:streamGetDict()

```

Page methods:

```

<boolean>      = <Page>:isOk()
<integer>      = <Page>:getNum()
<PDFRectangle> = <Page>:getMediaBox()
<PDFRectangle> = <Page>:getCropBox()
<boolean>      = <Page>:isCropped()
<number>       = <Page>:getMediaWidth()
<number>       = <Page>:getMediaHeight()
<number>       = <Page>:getCropWidth()
<number>       = <Page>:getCropHeight()
<PDFRectangle> = <Page>:getBleedBox()
<PDFRectangle> = <Page>:getTrimBox()
<PDFRectangle> = <Page>:getArtBox()
<integer>      = <Page>:getRotate()
<string>       = <Page>:getLastModified()
<Dict>         = <Page>:getBoxColorInfo()
<Dict>         = <Page>:getGroup()
<Stream>       = <Page>:getMetadata()
<Dict>         = <Page>:getPieceInfo()
<Dict>         = <Page>:getSeparationInfo()
<Dict>         = <Page>:getResourceDict()
<Object>       = <Page>:getAnnots()
<Links>        = <Page>:getLinks(<Catalog>)
<Object>       = <Page>:getContents()

```



PDFDoc methods:

```
<boolean> = <PDFDoc>:isOk()
<integer> = <PDFDoc>:getErrorCode()
<string> = <PDFDoc>:getErrorCodeName()
<string> = <PDFDoc>:getFileName()
<XRef> = <PDFDoc>:getXRef()
<Catalog> = <PDFDoc>:getCatalog()
<number> = <PDFDoc>:getPageMediaWidth()
<number> = <PDFDoc>:getPageMediaHeight()
<number> = <PDFDoc>:getPageCropWidth()
<number> = <PDFDoc>:getPageCropHeight()
<integer> = <PDFDoc>:getNumPages()
<string> = <PDFDoc>:readMetadata()
<Object> = <PDFDoc>:getStructTreeRoot()
<integer> = <PDFDoc>:findPage(<integer> object number, <integer> object genera-
tion)
<Links> = <PDFDoc>:getLinks(<integer>)
<LinkDest> = <PDFDoc>:findDest(<string>)
<boolean> = <PDFDoc>:isEncrypted()
<boolean> = <PDFDoc>:okToPrint()
<boolean> = <PDFDoc>:okToChange()
<boolean> = <PDFDoc>:okToCopy()
<boolean> = <PDFDoc>:okToAddNotes()
<boolean> = <PDFDoc>:isLinearized()
<Object> = <PDFDoc>:getDocInfo()
<Object> = <PDFDoc>:getDocInfoNF()
<integer> = <PDFDoc>:getPDFMajorVersion()
<integer> = <PDFDoc>:getPDFMinorVersion()
```

PDFRectangle methods:

```
<boolean> = <PDFRectangle>:isValid()
```

Stream methods:

```
<integer> = <Stream>:getKind()
<string> = <Stream>:getKindName()
          = <Stream>:reset()
          = <Stream>:close()
<integer> = <Stream>:getChar()
<integer> = <Stream>:lookChar()
<integer> = <Stream>:getRawChar()
<integer> = <Stream>:getUnfilteredChar()
          = <Stream>:unfilteredReset()
<integer> = <Stream>:getPos()
<boolean> = <Stream>:isBinary()
<Stream> = <Stream>:getUndecodedStream()
```



<Dict> = <Stream>:getDict()

StructElement methods:

```
<string>      = <StructElement>:getTypeName()
<integer>     = <StructElement>:getType()
<boolean>    = <StructElement>:isOk()
<boolean>    = <StructElement>:isBlock()
<boolean>    = <StructElement>:isInline()
<boolean>    = <StructElement>:isGrouping()
<boolean>    = <StructElement>:isContent()
<boolean>    = <StructElement>:isObjectRef()
<integer>    = <StructElement>:getMCID()
<Ref>        = <StructElement>:getObjectRef()
<Ref>        = <StructElement>:getParentRef()
<boolean>    = <StructElement>:hasPageRef()
<Ref>        = <StructElement>:getPageRef()
<StructTreeRoot> = <StructElement>:getStructTreeRoot()
<string>     = <StructElement>:getID()
<string>     = <StructElement>:getLanguage()
<integer>    = <StructElement>:getRevision()
              <StructElement>:setRevision(<unsigned integer>)
<string>     = <StructElement>:getTitle()
<string>     = <StructElement>:getExpandedAbbr()
<integer>    = <StructElement>:getNumChildren()
<StructElement> = <StructElement>:getChild()
              = <StructElement>:appendChild<StructElement>()
<integer>    = <StructElement>:getNumAttributes()
<Attribute> = <StructElement>:getAttribute(<integer>)
<string>    = <StructElement>:appendAttribute(<Attribute>)
<Attribute> = <StructElement>:findAttribute(<Attribute::Type>, boolean, At-
tribute::Owner)
<string>    = <StructElement>:getAltText()
<string>    = <StructElement>:getActualText()
<string>    = <StructElement>:getText(<boolean>)
<table>    = <StructElement>:getTextSpans()
```

StructTreeRoot methods:

```
<StructElement> = <StructTreeRoot>:findParentElement
<PDFDoc>       = <StructTreeRoot>:getDoc
<Dict>        = <StructTreeRoot>:getRoleMap
<Dict>        = <StructTreeRoot>:getClassMap
<integer>     = <StructTreeRoot>:getNumChildren
<StructElement> = <StructTreeRoot>:getChild
              <StructTreeRoot>:appendChild
<StructElement> = <StructTreeRoot>:findParentElement
```



TextSpan han only one method:

```
<string> = <TestSpan>:getText()
```

XRef methods:

```
<boolean> = <XRef>:isOk()  
<integer> = <XRef>:getErrorCode()  
<boolean> = <XRef>:isEncrypted()  
<boolean> = <XRef>:okToPrint()  
<boolean> = <XRef>:okToPrintHighRes()  
<boolean> = <XRef>:okToChange()  
<boolean> = <XRef>:okToCopy()  
<boolean> = <XRef>:okToAddNotes()  
<boolean> = <XRef>:okToFillForm()  
<boolean> = <XRef>:okToAccessibility()  
<boolean> = <XRef>:okToAssemble()  
<Object> = <XRef>:getCatalog()  
<Object> = <XRef>:fetch(<integer> object number, <integer> object generation)  
<Object> = <XRef>:getDocInfo()  
<Object> = <XRef>:getDocInfoNF()  
<integer> = <XRef>:getNumObjects()  
<integer> = <XRef>:getRootNum()  
<integer> = <XRef>:getRootGen()  
<integer> = <XRef>:getSize()  
<Object> = <XRef>:getTrailerDict()
```

There is an experimental function `epdf.openMemStream` that takes three arguments:

stream this is a (in low level Lua speak) light userdata object, i.e. a pointer to a sequence of bytes
length this is the length of the stream in bytes
name this is a unique identifier that us used for hashing the stream, so that multiple doesn't use more memory

Instead of a light userdata stream you can also pass a Lua string, in which case the given length is (at most) the string length.

The returned object can be used in the `img` library instead of a filename. 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.

7.3 The font library

The font library provides the interface into the internals of the font system, and also it contains helper functions to load traditional $\text{T}_{\text{E}}\text{X}$ font metrics formats. Other font loading functionality is provided by the `fontloader` library that will be discussed in the next section.



7.3.1 Loading a TFM file

The behavior documented in this subsection is considered stable in the sense that there will not be backward-incompatible changes any more.

```
<table> fnt = font.read_tfm(<string> name, <number> s)
```

The number 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 designsize of the font.

The internal structure of the metrics font table that is returned is explained in chapter 4.

7.3.2 Loading a VF file

The behavior documented in this subsection is considered stable in the sense that there will not be backward-incompatible changes any more.

```
<table> vf_fnt = font.read_vf(<string> name, <number> s)
```

The meaning of the number `s` and the format of the returned table are similar to the ones in the `read_tfm()` function.

7.3.3 The fonts array

The whole table of $\text{T}_{\text{E}}\text{X}$ fonts is accessible from Lua using a virtual array.

```
font.fonts[n] = { ... }  
<table> f = font.fonts[n]
```

See chapter 4 for the structure of the tables. Because this is a virtual array, you cannot call `pairs` on it, but see below for the `font.each` iterator.

The two metatable functions implementing the virtual array are:

```
<table> f = font.getfont(<number> n)  
font.setfont(<number> n, <table> f)
```

Note that at the moment, each access to the `font.fonts` or call to `font.getfont` creates a Lua table for the whole font. This process can be quite slow. In a later version of $\text{LuaT}_{\text{E}}\text{X}$, this interface will change (it will start using userdata objects instead of actual tables).

Also note the following: assignments can only be made to fonts that have already been defined in $\text{T}_{\text{E}}\text{X}$, but have not been accessed *at all* since that definition. This limits the usability of the write access to `font.fonts` quite a lot, a less stringent ruleset will likely be implemented later.

7.3.4 Checking a font’s status

You can test for the status of a font by calling this function:



```
<boolean> f = font.frozen(<number> n)
```

The return value is one of `true` (unassignable), `false` (can be changed) or `nil` (not a valid font at all).

7.3.5 Defining a font directly

You can define your own font into `font.fonts` by calling this function:

```
<number> i = font.define(<table> f)
```

The return value is the internal id number of the defined font (the index into `font.fonts`). If the font creation fails, an error is raised. The table is a font structure, as explained in chapter 4.

7.3.6 Projected next font id

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

This returns the font id number that would be returned by a `font.define` call if it was executed at this spot in the code flow. This is useful for virtual fonts that need to reference themselves.

7.3.7 Font id

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

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

7.3.8 Currently active font

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

This gets or sets the currently used font number.

7.3.9 Maximum font id

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

This is the largest used index in `font.fonts`.

7.3.10 Iterating over all fonts

```
for i,v in font.each() do  
    ...  
end
```



This is an iterator over each of the defined \TeX fonts. The first returned value is the index in `font.fonts`, the second the font itself, as a Lua table. The indices are listed incrementally, but they do not always form an array of consecutive numbers: in some cases there can be holes in the sequence.

7.4 The fontloader library

7.4.1 Getting quick information on a font

```
<table> info = fontloader.info(<string> filename)
```

This function returns either `nil`, or a `table`, or an array of small tables (in the case of a TrueType collection). The returned table(s) will contain some fairly interesting information items from the font(s) defined by the file:

key	type	explanation
fontname	string	the PostScript name of the font
fullname	string	the formal name of the font
familyname	string	the family name this font belongs to
weight	string	a string indicating the color value of the font
version	string	the internal font version
italicangle	float	the slant angle
units_per_em	number	1000 for PostScript-based fonts, usually 2048 for TrueType
pfminfo	table	(see section 7.4.5.1.6)

Getting information through this function is (sometimes much) more efficient than loading the font properly, and is therefore handy when you want to create a dictionary of available fonts based on a directory contents.

7.4.2 Loading an OPENTYPE or TRUETYPE file

If you want to use an OpenType font, you have to get the metric information from somewhere. Using the `fontloader` library, the simplest way to get that information is thus:

```
function load_font (filename)
  local metrics = nil
  local font = fontloader.open(filename)
  if font then
    metrics = fontloader.to_table(font)
    fontloader.close(font)
  end
  return metrics
end

myfont = load_font('/opt/tex/texmf/fonts/data/arial.ttf')
```

The main function call is



```
<userdata> f, <table> w = fontloader.open(<string> filename)
<userdata> f, <table> w = fontloader.open(<string> filename, <string> fontname)
```

The first return value is a userdata representation of the font. The second return value is a table containing any warnings and errors reported by fontloader while opening the font. In normal typesetting, you would probably ignore the second argument, but it can be useful for debugging purposes.

For TrueType collections (when filename ends in 'ttc') and dfont collections, you have to use a second string argument to specify which font you want from the collection. Use the `fontname` strings that are returned by `fontloader.info` for that.

To turn the font into a table, `fontloader.to_table` is used on the font returned by `fontloader.open`.

```
<table> f = fontloader.to_table(<userdata> font)
```

This table cannot be used directly by Lua_T_EX and should be turned into another one as described in chapter 4. Do not forget to store the `fontname` value in the `psname` field of the metrics table to be returned to Lua_T_EX, otherwise the font inclusion backend will not be able to find the correct font in the collection.

See section 7.4.5 for details on the userdata object returned by `fontloader.open()` and the layout of the `metrics` table returned by `fontloader.to_table()`.

The font file is parsed and partially interpreted by the font loading routines from FontForge. The file format can be OpenType, TrueType, TrueType Collection, cff, or Type1.

There are a few advantages to this approach compared to reading the actual font file ourselves:

- The font is automatically re-encoded, so that the `metrics` table for TrueType and OpenType fonts is using Unicode for the character indices.
- Many features are pre-processed into a format that is easier to handle than just the bare tables would be.
- PostScript-based OpenType fonts do not store the character height and depth in the font file, so the character boundingbox has to be calculated in some way.
- In the future, it may be interesting to allow Lua scripts access to the font program itself, perhaps even creating or changing the font.

A loaded font is discarded with:

```
fontloader.close(<userdata> font)
```

7.4.3 Applying a 'feature file'

You can apply a 'feature file' to a loaded font:

```
<table> errors = fontloader.apply_featurefile(<userdata> font, <string> filename)
```

A 'feature file' is a textual representation of the features in an OpenType font. See



http://www.adobe.com/devnet/opentype/afdko/topic_feature_file_syntax.html

and

<http://fontforge.sourceforge.net/featurefile.html>

for a more detailed description of feature files.

If the function fails, the return value is a table containing any errors reported by fontloader while applying the feature file. On success, `nil` is returned.

7.4.4 Applying an ‘AFM file’

You can apply an ‘afm file’ to a loaded font:

```
<table> errors = fontloader.apply_afmfile(<userdata> font, <string> filename)
```

An afm file is a textual representation of (some of) the meta information in a Type1 font. See

ftp://ftp.math.utah.edu/u/ma/hohn/linux/postscript/5004.AFM_Spec.pdf

for more information about afm files.

Note: If you `fontloader.open()` a Type1 file named `font.pfb`, the library will automatically search for and apply `font.afm` if it exists in the same directory as the file `font.pfb`. In that case, there is no need for an explicit call to `apply_afmfile()`.

If the function fails, the return value is a table containing any errors reported by fontloader while applying the AFM file. On success, `nil` is returned.

7.4.5 Fontloader font tables

As mentioned earlier, the return value of `fontloader.open()` is a userdata object. One way to have access to the actual metrics is to call `fontloader.to_table()` on this object, returning the table structure that is explained in the following subsections.

However, it turns out that the result from `fontloader.to_table()` sometimes needs very large amounts of memory (depending on the font’s complexity and size) so it is possible to access the userdata object directly.

- All top-level keys that would be returned by `to_table()` can also be accessed directly.
-
- The top-level key ‘glyphs’ returns a *virtual* array that allows indices from `f.glyphmin` to `(f.glyphmax)`.
- The items in that virtual array (the actual glyphs) are themselves also userdata objects, and each has accessors for all of the keys explained in the section ‘Glyph items’ below.
The top-level key ‘subfonts’ returns an *actual* array of userdata objects, one for each of the subfonts (or `nil`, if there are no subfonts).

A short example may be helpful. This code generates a printout of all the glyph names in the font `PunkNova.kern.otf`:




```

local f = fontloader.open('PunkNova.kern.otf')
print (f.fontname)
local i = 0
if f.glyphcnt > 0 then
  for i=f.glyphmin,f.glyphmax do
    local g = f.glyphs[i]
    if g then
      print(g.name)
    end
    i = i + 1
  end
end
fontloader.close(f)

```

In this case, the LuaTeX memory requirement stays below 100MB on the test computer, while the internal structure generated by `to_table()` needs more than 2GB of memory (the font itself is 6.9MB in disk size).

Only the top-level font, the subfont table entries, and the glyphs are virtual objects, everything else still produces normal Lua values and tables.

If you want to know the valid fields in a font or glyph structure, call the `fields` function on an object of a particular type (either glyph or font):

```

<table> fields = fontloader.fields(<userdata> font)
<table> fields = fontloader.fields(<userdata> font_glyph)

```

For instance:

```

local fields = fontloader.fields(f)
local fields = fontloader.fields(f.glyphs[0])

```

7.4.5.1 Table types

7.4.5.1.1 Top-level

The top-level keys in the returned table are (the explanations in this part of the documentation are not yet finished):

key	type	explanation
table_version	number	indicates the metrics version (currently 0.3)
fontname	string	PostScript font name
fullname	string	official (human-oriented) font name
familyname	string	family name
weight	string	weight indicator
copyright	string	copyright information
filename	string	the file name
version	string	font version



<code>italicangle</code>	float	slant angle
<code>units_per_em</code>	number	1000 for PostScript-based fonts, usually 2048 for TrueType
<code>ascent</code>	number	height of ascender in <code>units_per_em</code>
<code>descent</code>	number	depth of descender in <code>units_per_em</code>
<code>upos</code>	float	
<code>uwidth</code>	float	
<code>uniqueid</code>	number	
<code>glyphs</code>	array	
<code>glyphcnt</code>	number	number of included glyphs
<code>glyphmax</code>	number	maximum used index the glyphs array
<code>glyphmin</code>	number	minimum used index the glyphs array
<code>notdef_loc</code>	number	location of the <code>.notdef</code> glyph or <code>-1</code> when not present
<code>hasvmetrics</code>	number	
<code>onlybitmaps</code>	number	
<code>serifcheck</code>	number	
<code>isserif</code>	number	
<code>issans</code>	number	
<code>encodingchanged</code>	number	
<code>strokedfont</code>	number	
<code>use_typo_metrics</code>	number	
<code>weight_width_slope_only</code>	number	
<code>head_optimized_for_cleartype</code>	number	
<code>uni_interp</code>	enum	<code>unset</code> , <code>none</code> , <code>adobe</code> , <code>greek</code> , <code>japanese</code> , <code>trad_chinese</code> , <code>simp_chinese</code> , <code>korean</code> , <code>ams</code>
<code>origname</code>	string	the file name, as supplied by the user
<code>map</code>	table	
<code>private</code>	table	
<code>xuid</code>	string	
<code>pfminfo</code>	table	
<code>names</code>	table	
<code>cidinfo</code>	table	
<code>subfonts</code>	array	
<code>comments</code>	string	
<code>fontlog</code>	string	
<code>cvt_names</code>	string	
<code>anchor_classes</code>	table	
<code>ttf_tables</code>	table	
<code>ttf_tab_saved</code>	table	
<code>kerns</code>	table	
<code>vkerns</code>	table	
<code>texdata</code>	table	
<code>lookups</code>	table	
<code>gpos</code>	table	
<code>gsub</code>	table	
<code>mm</code>	table	



chosename	string
macstyle	number
fondname	string
fontstyle_id	number
fontstyle_name	table
strokewidth	float
mark_classes	table
creationtime	number
modificationtime	number
os2_version	number
sfd_version	number
math	table
validation_state	table
horiz_base	table
vert_base	table
extrema_bound	number
truetype	boolean signals a TrueType font

7.4.5.1.2 Glyph items

The `glyphs` is an array containing the per-character information (quite a few of these are only present if nonzero).

key	type	explanation
name	string	the glyph name
unicode	number	unicode code point, or -1
boundingbox	array	array of four numbers, see note below
width	number	only for horizontal fonts
vwidth	number	only for vertical fonts
tsidebearing	number	only for vertical ttf/otf fonts, and only if nonzero
lsidebearing	number	only if nonzero and not equal to boundingbox[1]
class	string	one of "none", "base", "ligature", "mark", "component" (if not present, the glyph class is 'automatic')
kerns	array	only for horizontal fonts, if set
vkerns	array	only for vertical fonts, if set
dependents	array	linear array of glyph name strings, only if nonempty
lookups	table	only if nonempty
ligatures	table	only if nonempty
anchors	table	only if set
comment	string	only if set
tex_height	number	only if set
tex_depth	number	only if set
italic_correction	number	only if set
top_accent	number	only if set
is_extended_shape	number	only if this character is part of a math extension list
altuni	table	alternate Unicode items
vert_variants	table	



horiz_variants	table
mathkern	table

On [boundingbox](#): The boundingbox information for TrueType fonts and TrueType-based otf fonts is read directly from the font file. PostScript-based fonts do not have this information, so the boundingbox of traditional PostScript fonts is generated by interpreting the actual bezier curves to find the exact boundingbox. This can be a slow process, so the boundingboxes of PostScript-based otf fonts (and raw cff fonts) are calculated using an approximation of the glyph shape based on the actual glyph points only, instead of taking the whole curve into account. This means that glyphs that have missing points at extrema will have a too-tight boundingbox, but the processing is so much faster that in our opinion the tradeoff is worth it.

The [kerns](#) and [vkerns](#) are linear arrays of small hashes:

key	type	explanation
char	string	
off	number	
lookup	string	

The [lookups](#) is a hash, based on lookup subtable names, with the value of each key inside that a linear array of small hashes:

key	type	explanation
type	enum	position , pair , substitution , alternate , multiple , ligature , lcaret , kerning , vkerning , anchors , contextpos , contextsub , chainpos , chain-sub , reversesub , max , kernback , vkernback
specification	table	extra data

For the first seven values of [type](#), there can be additional sub-information, stored in the sub-table [specification](#):

value	type	explanation
position	table	a table of the offset_specs type
pair	table	one string: paired , and an array of one or two offset_specs tables: offsets
substitution	table	one string: variant
alternate	table	one string: components
multiple	table	one string: components
ligature	table	two strings: components , char
lcaret	array	linear array of numbers

Tables for [offset_specs](#) contain up to four number-valued fields: [x](#) (a horizontal offset), [y](#) (a vertical offset), [h](#) (an advance width correction) and [v](#) (an advance height correction).

The [ligatures](#) is a linear array of small hashes:

key	type	explanation
lig	table	uses the same substructure as a single item in the lookups table explained above
char	string	



components	array	linear array of named components
ccnt	number	

The `anchor` table is indexed by a string signifying the anchor type, which is one of

key	type	explanation
mark	table	placement mark
basechar	table	mark for attaching combining items to a base char
baselig	table	mark for attaching combining items to a ligature
basemark	table	generic mark for attaching combining items to connect to
centry	table	cursive entry point
cexit	table	cursive exit point

The content of these is a short array of defined anchors, with the entry keys being the anchor names. For all except `baselig`, the value is a single table with this definition:

key	type	explanation
x	number	x location
y	number	y location
ttf_pt_index	number	truetype point index, only if given

For `baselig`, the value is a small array of such anchor sets sets, one for each constituent item of the ligature.

For clarification, an anchor table could for example look like this :

```
['anchor'] = {
  ['basemark'] = {
    ['Anchor-7'] = { ['x']=170, ['y']=1080 }
  },
  ['mark'] = {
    ['Anchor-1'] = { ['x']=160, ['y']=810 },
    ['Anchor-4'] = { ['x']=160, ['y']=800 }
  },
  ['baselig'] = {
    [1] = { ['Anchor-2'] = { ['x']=160, ['y']=650 } },
    [2] = { ['Anchor-2'] = { ['x']=460, ['y']=640 } }
  }
}
```

Note: The `baselig` table can be sparse!

7.4.5.1.3 map table

The top-level map is a list of encoding mappings. Each of those is a table itself.

key	type	explanation
enccount	number	
encmax	number	



backmax	number	
remap	table	
map	array	non-linear array of mappings
backmap	array	non-linear array of backward mappings
enc	table	

The `remap` table is very small:

key	type	explanation
firstenc	number	
lastenc	number	
infont	number	

The `enc` table is a bit more verbose:

key	type	explanation
enc_name	string	
char_cnt	number	
char_max	number	
unicode	array	of Unicode position numbers
psnames	array	of PostScript glyph names
builtin	number	
hidden	number	
only_1byte	number	
has_1byte	number	
has_2byte	number	
is_unicodebmp	number	only if nonzero
is_unicodedefull	number	only if nonzero
is_custom	number	only if nonzero
is_original	number	only if nonzero
is_compact	number	only if nonzero
is_japanese	number	only if nonzero
is_korean	number	only if nonzero
is_tradchinese	number	only if nonzero [name?]
is_simplechinese	number	only if nonzero
low_page	number	
high_page	number	
iconv_name	string	
iso_2022_escape	string	

7.4.5.1.4 private table

This is the font's private PostScript dictionary, if any. Keys and values are both strings.

7.4.5.1.5 cidinfo table

key	type	explanation
registry	string	



ordering	string
supplement	number
version	number

7.4.5.1.6 pfminfo table

The `pfminfo` table contains most of the OS/2 information:

key	type	explanation
pfmset	number	
winascent_add	number	
windescent_add	number	
hheadascent_add	number	
hheaddescent_add	number	
typoascent_add	number	
typodescent_add	number	
subsuper_set	number	
panose_set	number	
hheadset	number	
vheadset	number	
pfmfamily	number	
weight	number	
width	number	
avgwidth	number	
firstchar	number	
lastchar	number	
fstype	number	
linegap	number	
vlinegap	number	
hhead_ascent	number	
hhead_descent	number	
os2_typoascent	number	
os2_typodescent	number	
os2_typolinegap	number	
os2_winascent	number	
os2_windescent	number	
os2_subxsize	number	
os2_subysize	number	
os2_subxoff	number	
os2_subyoff	number	
os2_supxsize	number	
os2_supysize	number	
os2_supxoff	number	
os2_supyoff	number	
os2_strikeysize	number	
os2_strikeypos	number	
os2_family_class	number	



os2_xheight	number	
os2_capheight	number	
os2_defaultchar	number	
os2_breakchar	number	
os2_vendor	string	
codepages	table	A two-number array of encoded code pages
unicoderanges	table	A four-number array of encoded unicode ranges
panose	table	

The `panose` subtable has exactly 10 string keys:

key	type	explanation
familytype	string	Values as in the OpenType font specification: Any , No Fit , Text and Display , Script , Decorative , Pictorial
serifstyle	string	See the OpenType font specification for values
weight	string	id.
proportion	string	id.
contrast	string	id.
strokevariation	string	id.
armstyle	string	id.
letterform	string	id.
midline	string	id.
xheight	string	id.

7.4.5.1.7 names table

Each item has two top-level keys:

key	type	explanation
lang	string	language for this entry
names	table	

The `names` keys are the actual TrueType name strings. The possible keys are:

key	explanation
copyright	
family	
subfamily	
uniqueid	
fullname	
version	
postscriptname	
trademark	
manufacturer	
designer	
descriptor	
venderurl	
designerurl	



license
 licenseurl
 idontknow
 preffamilyname
 prefmodifiers
 compatfull
 sampletext
 cidfindfontname
 wwsfamily
 wwssubfamily

7.4.5.1.8 anchor_classes table

The anchor_classes classes:

key	type	explanation
name	string	a descriptive id of this anchor class
lookup	string	
type	string	one of <code>mark</code> , <code>mkmk</code> , <code>curs</code> , <code>mklg</code>

7.4.5.1.9 gpos table

The `gpos` table has one array entry for each lookup. (The `gpos_` prefix is somewhat redundant.)

key	type	explanation
type	string	one of <code>gpos_single</code> , <code>gpos_pair</code> , <code>gpos_cursive</code> , <code>gpos_mark2base</code> , <code>gpos_mark2ligature</code> , <code>gpos_mark2mark</code> , <code>gpos_context</code> , <code>gpos_contextchain</code>
flags	table	
name	string	
features	array	
subtables	array	

The flags table has a true value for each of the lookup flags that is actually set:

key	type	explanation
r2l	boolean	
ignorebaseglyphs	boolean	
ignoreligatures	boolean	
ignorerecombiningmarks	boolean	
mark_class	string	

The features subtable items of `gpos` have:

key	type	explanation
tag	string	
scripts	table	



The scripts table within features has:

key	type	explanation
script	string	
langs	array of strings	

The subtables table has:

key	type	explanation
name	string	
suffix	string	(only if used)
anchor_classes	number	(only if used)
vertical_kerning	number	(only if used)
kernclass	table	(only if used)

The kernclass with subtables table has:

key	type	explanation
firsts	array of strings	
seconds	array of strings	
lookup	string or array	associated lookup(s)
offsets	array of numbers	

Note: the kernclass (as far as we can see) always has one entry so it could be one level deep instead. Also the seconds start at [2] which is close to the fontforge internals so we keep that too.

7.4.5.1.10 gsub table

This has identical layout to the gpos table, except for the type:

key	type	explanation
type	string	one of <code>gsub_single</code> , <code>gsub_multiple</code> , <code>gsub_alternate</code> , <code>gsub_ligature</code> , <code>gsub_context</code> , <code>gsub_contextchain</code> , <code>gsub_reversecontextchain</code>

7.4.5.1.11 ttf_tables and ttf_tab_saved tables

key	type	explanation
tag	string	
len	number	
maxlen	number	
data	number	

7.4.5.1.12 mm table

key	type	explanation
axes	table	array of axis names
instance_count	number	



positions	table	array of instance positions (#axes * instances)
defweights	table	array of default weights for instances
cdv	string	
ndv	string	
axismaps	table	

The `axismaps`:

key	type	explanation
blends	table	an array of blend points
designs	table	an array of design values
min	number	
def	number	
max	number	

7.4.5.1.13 `mark_classes` table

The keys in this table are mark class names, and the values are a space-separated string of glyph names in this class.

7.4.5.1.14 `math` table

ScriptPercentScaleDown
ScriptScriptPercentScaleDown
DelimitedSubFormulaMinHeight
DisplayOperatorMinHeight
MathLeading
AxisHeight
AccentBaseHeight
FlattenedAccentBaseHeight
SubscriptShiftDown
SubscriptTopMax
SubscriptBaselineDropMin
SuperscriptShiftUp
SuperscriptShiftUpCramped
SuperscriptBottomMin
SuperscriptBaselineDropMax
SubSuperscriptGapMin
SuperscriptBottomMaxWithSubscript
SpaceAfterScript
UpperLimitGapMin
UpperLimitBaselineRiseMin
LowerLimitGapMin
LowerLimitBaselineDropMin
StackTopShiftUp
StackTopDisplayStyleShiftUp
StackBottomShiftDown



StackBottomDisplayStyleShiftDown
StackGapMin
StackDisplayStyleGapMin
StretchStackTopShiftUp
StretchStackBottomShiftDown
StretchStackGapAboveMin
StretchStackGapBelowMin
FractionNumeratorShiftUp
FractionNumeratorDisplayStyleShiftUp
FractionDenominatorShiftDown
FractionDenominatorDisplayStyleShiftDown
FractionNumeratorGapMin
FractionNumeratorDisplayStyleGapMin
FractionRuleThickness
FractionDenominatorGapMin
FractionDenominatorDisplayStyleGapMin
SkewedFractionHorizontalGap
SkewedFractionVerticalGap
OverbarVerticalGap
OverbarRuleThickness
OverbarExtraAscender
UnderbarVerticalGap
UnderbarRuleThickness
UnderbarExtraDescender
RadicalVerticalGap
RadicalDisplayStyleVerticalGap
RadicalRuleThickness
RadicalExtraAscender
RadicalKernBeforeDegree
RadicalKernAfterDegree
RadicalDegreeBottomRaisePercent
MinConnectorOverlap
FractionDelimiterSize
FractionDelimiterDisplayStyleSize

7.4.5.1.15 validation_state table

key	explanation
bad_ps_fontname	
bad_glyph_table	
bad_cff_table	
bad_metrics_table	
bad_cmap_table	
bad_bitmaps_table	
bad_gx_table	
bad_ot_table	



bad_os2_version
bad_sfnt_header

7.4.5.1.16 **horiz_base and vert_base table**

key	type	explanation
tags	table	an array of script list tags
scripts	table	

The `scripts` subtable:

key	type	explanation
baseline	table	
default_baseline	number	
lang	table	

The `lang` subtable:

key	type	explanation
tag	string	a script tag
ascent	number	
descent	number	
features	table	

The `features` points to an array of tables with the same layout except that in those nested tables, the tag represents a language.

7.4.5.1.17 **altuni table**

An array of alternate Unicode values. Inside that array are hashes with:

key	type	explanation
unicode	number	this glyph is also used for this unicode
variant	number	the alternative is driven by this unicode selector

7.4.5.1.18 **vert_variants and horiz_variants table**

key	type	explanation
variants	string	
italic_correction	number	
parts	table	

The `parts` table is an array of smaller tables:

key	type	explanation
component	string	
extender	number	
start	number	



end	number
advance	number

7.4.5.1.19 mathkern table

key	type	explanation
top_right	table	
bottom_right	table	
top_left	table	
bottom_left	table	

Each of the subtables is an array of small hashes with two keys:

key	type	explanation
height	number	
kern	number	

7.4.5.1.20 kerns table

Substructure is identical to the per-glyph subtable.

7.4.5.1.21 vkerns table

Substructure is identical to the per-glyph subtable.

7.4.5.1.22 texdata table

key	type	explanation
type	string	unset, text, math, mathext
params	array	22 font numeric parameters

7.4.5.1.23 lookups table

Top-level `lookups` is quite different from the ones at character level. The keys in this hash are strings, the values the actual lookups, represented as dictionary tables.

key	type	explanation
type	string	
format	enum	one of <code>glyphs</code> , <code>class</code> , <code>coverage</code> , <code>reversecoverage</code>
tag	string	
current_class	array	
before_class	array	
after_class	array	
rules	array	an array of rule items

Rule items have one common item and one specialized item:



key	type	explanation
lookups	array	a linear array of lookup names
glyphs	array	only if the parent's format is glyphs
class	array	only if the parent's format is class
coverage	array	only if the parent's format is coverage
reversecoverage	array	only if the parent's format is reversecoverage

A glyph table is:

key	type	explanation
names	string	
back	string	
fore	string	

A class table is:

key	type	explanation
current	array	of numbers
before	array	of numbers
after	array	of numbers

coverage:

key	type	explanation
current	array	of strings
before	array	of strings
after	array	of strings

reversecoverage:

key	type	explanation
current	array	of strings
before	array	of strings
after	array	of strings
replacements	string	

7.5 The `img` library

The `img` library can be used as an alternative to `\pdfximage` and `\pdfrefximage`, and the associated 'satellite' commands like `\pdfximagebbox`. Image objects can also be used within virtual fonts via the `image` command listed in section 4.2.

7.5.1 `img.new`

```
<image> var = img.new()
<image> var = img.new(<table> image_spec)
```



This function creates a userdata object of type 'image'. The `image_spec` argument is optional. If it is given, it must be a table, and that table must contain a `filename` key. A number of other keys can also be useful, these are explained below.

You can either say

```
a = img.new()
```

followed by

```
a.filename = "foo.png"
```

or you can put the file name (and some or all of the other keys) into a table directly, like so:

```
a = img.new({filename='foo.pdf', page=1})
```

The generated `<image>` userdata object allows access to a set of user-specified values as well as a set of values that are normally filled in and updated automatically by LuaTeX itself. Some of those are derived from the actual image file, others are updated to reflect the pdf output status of the object.

There is one required user-specified field: the file name (`filename`). It can optionally be augmented by the requested image dimensions (`width`, `depth`, `height`), user-specified image attributes (`attr`), the requested pdf page identifier (`page`), the requested boundingbox (`pagebox`) for pdf inclusion, the requested color space object (`colorspace`).

The function `img.new` does not access the actual image file, it just creates the `<image>` userdata object and initializes some memory structures. The `<image>` object and its internal structures are automatically garbage collected.

Once the image is scanned, all the values in the `<image>` except `width`, `height` and `depth`, become frozen, and you cannot change them any more.

You can use `pdf.setignoreunknownimages(1)` (or at the TeX end the `\pdfvariable ignoreunknownimages`) to get around a quit when no known image type is found (based on name or preamble). Beware: this will not catch invalid images and we cannot guarantee side effects. A zero dimension image is still included when requested. No special flags are set. A proper workflow will not rely in such a catch but make sure that images are valid.

7.5.2 `img.keys`

```
<table> keys = img.keys()
```

This function returns a list of all the possible `image_spec` keys, both user-supplied and automatic ones.

field name	type	description
<code>attr</code>	string	the image attributes for LuaTeX
<code>bbox</code>	table	table with 4 boundingbox dimensions <code>llx</code> , <code>lly</code> , <code>urx</code> , and <code>ury</code> overruling the <code>pagebox</code> entry
<code>colordepth</code>	number	the number of bits used by the color space



colorspace	number	the color space object number
depth	number	the image depth for LuaTeX (in scaled points)
filename	string	the image file name
filepath	string	the full (expanded) file name of the image
height	number	the image height for LuaTeX (in scaled points)
imagetype	string	one of <code>pdf</code> , <code>png</code> , <code>jpg</code> , <code>jp2</code> , <code>jbig2</code> , or <code>nil</code>
index	number	the pdf image name suffix
objnum	number	the pdf image object number
page	??	the identifier for the requested image page (type is number or string, default is the number 1)
pagebox	string	the requested bounding box, one of <code>none</code> , <code>media</code> , <code>crop</code> , <code>bleed</code> , <code>trim</code> , <code>art</code>
pages	number	the total number of available pages
rotation	number	the image rotation from included pdf file, in multiples of 90 deg.
stream	string	the raw stream data for an <code>/XObject /Form</code> object
transform	number	the image transform, integer number 0..7
width	number	the image width for LuaTeX (in scaled points)
xres	number	the horizontal natural image resolution (in dpi)
xsize	number	the natural image width
yres	number	the vertical natural image resolution (in dpi)
ysize	number	the natural image height
visiblefileame	string	when set, this name will find its way in the pdf file as <code>PTEX</code> specification; when an empty string is assigned nothing is written to file, otherwise the natural filename is taken

A running (undefined) dimension in `width`, `height`, or `depth` is represented as `nil` in Lua, so if you want to load an image at its ‘natural’ size, you do not have to specify any of those three fields.

The `stream` parameter allows to fabricate an `/XObject /Form` object from a string giving the stream contents, e.g., for a filled rectangle:

```
a.stream = "0 0 20 10 re f"
```

When writing the image, an `/XObject /Form` object is created, like with embedded pdf file writing. The object is written out only once. The `stream` key requires that also the `bbox` table is given. The `stream` key conflicts with the `filename` key. The `transform` key works as usual also with `stream`.

The `bbox` key needs a table with four boundingbox values, e.g.:

```
a.bbox = {"30bp", 0, "225bp", "200bp"}
```

This replaces and overrules any given `pagebox` value; with given `bbox` the box dimensions coming with an embedded pdf file are ignored. The `xsize` and `ysize` dimensions are set accordingly, when the image is scaled. The `bbox` parameter is ignored for non-pdf images.

The `transform` allows to mirror and rotate the image in steps of 90 deg. The default value 0 gives an unmirrored, unrotated image. Values 1 – 3 give counterclockwise rotation by 90, 180,



or 270 degrees, whereas with values 4 – 7 the image is first mirrored and then rotated counterclockwise by 90, 180, or 270 degrees. The `transform` operation gives the same visual result as if you would externally preprocess the image by a graphics tool and then use it by LuaTeX. If a pdf file to be embedded already contains a `/Rotate` specification, the rotation result is the combination of the `/Rotate` rotation followed by the `transform` operation.

7.5.3 `img.scan`

```
<image> var = img.scan(<image> var)
<image> var = img.scan(<table> image_spec)
```

When you say `img.scan(a)` for a new image, the file is scanned, and variables such as `xsize`, `ysize`, image `type`, number of `pages`, and the resolution are extracted. Each of the `width`, `height`, `depth` fields are set up according to the image dimensions, if they were not given an explicit value already. An image file will never be scanned more than once for a given image variable. With all subsequent `img.scan(a)` calls only the dimensions are again set up (if they have been changed by the user in the meantime).

For ease of use, you can do right-away a

```
<image> a = img.scan ({ filename = "foo.png" })
```

without a prior `img.new`.

Nothing is written yet at this point, so you can do `a=img.scan`, retrieve the available info like image width and height, and then throw away `a` again by saying `a=nil`. In that case no image object will be reserved in the PDF, and the used memory will be cleaned up automatically.

7.5.4 `img.copy`

```
<image> var = img.copy(<image> var)
<image> var = img.copy(<table> image_spec)
```

If you say `a = b`, then both variables point to the same `<image>` object. if you want to write out an image with different sizes, you can do a `b=img.copy(a)`.

Afterwards, `a` and `b` still reference the same actual image dictionary, but the dimensions for `b` can now be changed from their initial values that were just copies from `a`.

7.5.5 `img.write`

```
<image> var = img.write(<image> var)
<image> var = img.write(<table> image_spec)
```

By `img.write(a)` a pdf object number is allocated, and a whatsit node of subtype `pdf_refximage` is generated and put into the output list. By this the image `a` is placed into the page stream, and the image file is written out into an image stream object after the shipping of the current page is finished.



Again you can do a terse call like

```
img.write ({ filename = "foo.png" })
```

The `<image>` variable is returned in case you want it for later processing.

7.5.6 `img.immediatewrite`

```
<image> var = img.immediatewrite(<image> var)
<image> var = img.immediatewrite(<table> image_spec)
```

By `img.immediatewrite(a)` a pdf object number is allocated, and the image file for image `a` is written out immediately into the pdf file as an image stream object (like with `\immediate\pdfx-image`). The object number of the image stream dictionary is then available by the `objnum` key. No `pdf_refximage` whatsit node is generated. You will need an `img.write(a)` or `img.node(a)` call to let the image appear on the page, or reference it by another trick; else you will have a dangling image object in the pdf file.

Also here you can do a terse call like

```
a = img.immediatewrite ({ filename = "foo.png" })
```

The `<image>` variable is returned and you will most likely need it.

7.5.7 `img.node`

```
<node> n = img.node(<image> var)
<node> n = img.node(<table> image_spec)
```

This function allocates a pdf object number and returns a whatsit node of subtype `pdf_refximage`, filled with the image parameters `width`, `height`, `depth`, and `objnum`. Also here you can do a terse call like:

```
n = img.node ({ filename = "foo.png" })
```

This example outputs an image:

```
node.write(img.node{filename="foo.png"})
```

7.5.8 `img.types`

```
<table> types = img.types()
```

This function returns a list with the supported image file type names, currently these are `pdf`, `png`, `jpg`, `jp2` (JPEG 2000), and `jbig2`.

7.5.9 `img_boxes`

```
<table> boxes = img_boxes()
```



This function returns a list with the supported pdf page box names, currently these are `media`, `crop`, `bleed`, `trim`, and `art` (all in lowercase letters).

7.6 The `kpse` library

This library provides two separate, but nearly identical interfaces to the `kpathsea` file search functionality: there is a ‘normal’ procedural interface that shares its `kpathsea` instance with Lua \TeX itself, and an object oriented interface that is completely on its own.

7.6.1 `kpse.set_program_name` and `kpse.new`

Before the search library can be used at all, its database has to be initialized. There are three possibilities, two of which belong to the procedural interface.

First, when Lua \TeX is used to typeset documents, this initialization happens automatically and the `kpathsea` executable and program names are set to `luatex` (that is, unless explicitly prohibited by the user’s startup script. See section 2.1 for more details).

Second, in \TeX Lua mode, the initialization has to be done explicitly via the `kpse.set_program_name` function, which sets the `kpathsea` executable (and optionally program) name.

```
kpse.set_program_name(<string> name)
kpse.set_program_name(<string> name, <string> progname)
```

The second argument controls the use of the ‘dotted’ values in the `texmf.cnf` configuration file, and defaults to the first argument.

Third, if you prefer the object oriented interface, you have to call a different function. It has the same arguments, but it returns a userdata variable.

```
local kpathsea = kpse.new(<string> name)
local kpathsea = kpse.new(<string> name, <string> progname)
```

Apart from these two functions, the calling conventions of the interfaces are identical. Depending on the chosen interface, you either call `kpse.find_file()` or `kpathsea:find_file()`, with identical arguments and return values.

7.6.2 `find_file`

The most often used function in the library is `find_file`:

```
<string> f = kpse.find_file(<string> filename)
<string> f = kpse.find_file(<string> filename, <string> ftype)
<string> f = kpse.find_file(<string> filename, <boolean> mustexist)
<string> f = kpse.find_file(<string> filename, <string> ftype, <boolean> mustexist)
<string> f = kpse.find_file(<string> filename, <string> ftype, <number> dpi)
```



Arguments:

filename

the name of the file you want to find, with or without extension.

ftype

maps to the `-format` argument of `kpsewhich`. The supported `ftype` values are the same as the ones supported by the standalone `kpsewhich` program:

<code>gf</code>	TeX system sources
<code>pk</code>	PostScript header
<code>bitmap font</code>	Troff fonts
<code>tfm</code>	type1 fonts
<code>afm</code>	vf
<code>base</code>	dvips config
<code>bib</code>	ist
<code>bst</code>	truetype fonts
<code>cnf</code>	type42 fonts
<code>ls-R</code>	web2c files
<code>fmt</code>	other text files
<code>map</code>	other binary files
<code>mem</code>	misc fonts
<code>mf</code>	web
<code>mfpool</code>	cweb
<code>mft</code>	enc files
<code>mp</code>	cmap files
<code>mppool</code>	subfont definition files
MetaPost support	opentype fonts
<code>ocp</code>	pdftex config
<code>ofm</code>	lig files
<code>opl</code>	texmfscripts
<code>otp</code>	lua
<code>ovf</code>	font feature files
<code>ovp</code>	cid maps
<code>graphic/figure</code>	mlbib
<code>tex</code>	mlbst
TeX system documentation	clua
<code>texpool</code>	

The default type is `tex`. Note: this is different from `kpsewhich`, which tries to deduce the file type itself from looking at the supplied extension.

`mustexist`

is similar to `kpsewhich`'s `-must-exist`, and the default is `false`. If you specify `true` (or a non-zero integer), then the `kpse` library will search the disk as well as the `ls-R` databases.

`dpi`

This is used for the size argument of the formats `pk`, `gf`, and `bitmap font`.

7.6.3 Lookup

A more powerful (but slower) generic method for finding files is also available. It returns a string



for each found file.

```
<string> f, ... = kpse.lookup(<string> filename, <table> options)
```

The options match commandline arguments from `kpsewhich`:

key	type	description
debug	number	set debugging flags for this lookup
format	string	use specific file type (see list above)
dpi	number	use this resolution for this lookup; default 600
path	string	search in the given path
all	boolean	output all matches, not just the first
mustexist	boolean	search the disk as well as ls-R if necessary
mktexpk	boolean	disable/enable mktexpk generation for this lookup
mktetex	boolean	disable/enable mktetex generation for this lookup
mktexmf	boolean	disable/enable mktexmf generation for this lookup
mktexfm	boolean	disable/enable mktexfm generation for this lookup
subdir	string or table	only output matches whose directory part ends with the given string(s)

7.6.4 `init_prog`

Extra initialization for programs that need to generate bitmap fonts.

```
kpse.init_prog(<string> prefix, <number> base_dpi, <string> mfmode)
kpse.init_prog(<string> prefix, <number> base_dpi, <string> mfmode, <string>
fallback)
```

7.6.5 `readable_file`

Test if an (absolute) file name is a readable file.

```
<string> f = kpse.readable_file(<string> name)
```

The return value is the actual absolute filename you should use, because the disk name is not always the same as the requested name, due to aliases and system-specific handling under e.g. msdos. Returns `nil` if the file does not exist or is not readable.

7.6.6 `expand_path`

Like `kpsewhich`'s `-expand-path`:

```
<string> r = kpse.expand_path(<string> s)
```

7.6.7 `expand_var`

Like `kpsewhich`'s `-expand-var`:

```
<string> r = kpse.expand_var(<string> s)
```



7.6.8 expand_braces

Like kpsewhich's `-expand-braces`:

```
<string> r = kpse.expand_braces(<string> s)
```

7.6.9 show_path

Like kpsewhich's `-show-path`:

```
<string> r = kpse.show_path(<string> ftype)
```

7.6.10 var_value

Like kpsewhich's `-var-value`:

```
<string> r = kpse.var_value(<string> s)
```

7.6.11 version

Returns the kpathsea version string.

```
<string> r = kpse.version()
```

7.7 The lang library

This library provides the interface to Lua_{TEX}'s structure representing a language, and the associated functions.

```
<language> l = lang.new()  
<language> l = lang.new(<number> id)
```

This function creates a new userdata object. An object of type `<language>` is the first argument to most of the other functions in the `lang` library. These functions can also be used as if they were object methods, using the colon syntax.

Without an argument, the next available internal id number will be assigned to this object. With argument, an object will be created that links to the internal language with that id number.

```
<number> n = lang.id(<language> l)
```

returns the internal `\language` id number this object refers to.

```
<string> n = lang.hyphenation(<language> l)  
lang.hyphenation(<language> l, <string> n)
```

Either returns the current hyphenation exceptions for this language, or adds new ones. The syntax of the string is explained in section 3.3.



```
lang.clear_hyphenation(<language> l)
```

Clears the exception dictionary (string) for this language.

```
<string> n = lang.clean(<language> l, <string> o)
<string> n = lang.clean(<string> o)
```

Creates a hyphenation key from the supplied hyphenation value. The syntax of the argument string is explained in section 3.3. This function is useful if you want to do something else based on the words in a dictionary file, like spell-checking.

```
<string> n = lang.patterns(<language> l)
lang.patterns(<language> l, <string> n)
```

Adds additional patterns for this language object, or returns the current set. The syntax of this string is explained in section 3.3.

```
lang.clear_patterns(<language> l)
```

Clears the pattern dictionary for this language.

```
<number> n = lang.prehyphenchar(<language> l)
lang.prehyphenchar(<language> l, <number> n)
```

Gets or sets the ‘pre-break’ hyphen character for implicit hyphenation in this language (initially the hyphen, decimal 45).

```
<number> n = lang.posthyphenchar(<language> l)
lang.posthyphenchar(<language> l, <number> n)
```

Gets or sets the ‘post-break’ hyphen character for implicit hyphenation in this language (initially null, decimal 0, indicating emptiness).

```
<number> n = lang.preexhyphenchar(<language> l)
lang.preexhyphenchar(<language> l, <number> n)
```

Gets or sets the ‘pre-break’ hyphen character for explicit hyphenation in this language (initially null, decimal 0, indicating emptiness).

```
<number> n = lang.postexhyphenchar(<language> l)
lang.postexhyphenchar(<language> l, <number> n)
```

Gets or sets the ‘post-break’ hyphen character for explicit hyphenation in this language (initially null, decimal 0, indicating emptiness).

```
<boolean> success = lang.hyphenate(<node> head)
<boolean> success = lang.hyphenate(<node> head, <node> tail)
```

Inserts hyphenation points (discretionary nodes) in a node list. If `tail` is given as argument, processing stops on that node. Currently, `success` is always true if `head` (and `tail`, if specified) are proper nodes, regardless of possible other errors.



Hyphenation works only on ‘characters’, a special subtype of all the glyph nodes with the node subtype having the value 1. Glyph nodes with different subtypes are not processed. See section 3.1 for more details.

The following two commands can be used to set or query hj codes:

```
lang.sethjcode(<language> l, <number> char, <number> usedchar)
<number> usedchar = lang.gethjcode(<language> l, <number> char)
```

When you set a hjcode the current sets get initialized unless the set was already initialized due to `\savingshyphcodes` being larger than zero.

7.8 The lua library

This library contains one read-only item:

```
<string> s = lua.version
```

This returns the Lua version identifier string. The value is currently Lua 5.1.

7.8.1 LUA bytecode registers

Lua registers can be used to communicate Lua functions across Lua 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.

Note: The function must not contain any upvalues. Currently, functions containing upvalues can be stored (and their upvalues are set to `nil`), but this is an artifact of the current Lua implementation and thus subject to change.

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 `iniTEX`. 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 `iniTeX`.

7.8.2 LUA 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.

7.9 The `mplib` library

The MetaPost library interface registers itself in the table `mplib`. It is based on `mplib` version 1.9991.

7.9.1 `mplib.new`

To create a new MetaPost instance, call

```
<mpinstance> mp = mp-lib.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>interaction</code>	string	the interaction mode, one of <code>batch</code> , <code>nonstop</code> , <code>scroll</code> , <code>errorstop</code>	<code>errorstop</code>
<code>job_name</code>	string	<code>--jobname</code>	<code>mpout</code>
<code>find_file</code>	function	a function to find files	only local files

The `find_file` function should be of this form:

```
<string> found = finder (<string> name, <string> mode, <string> type)
```

with:

name the requested file

mode the file mode: `r` or `w`

type 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 `mp:execute()` call.

7.9.2 `mp: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:

<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.

7.9.3 `mp:execute`

You can ask the MetaPost interpreter to run a chunk of code by calling

```
<table> rettable = mp:execute('metapost language chunk')
```

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.

7.9.4 `mp:finish`

```
<table> rettable = mp:finish()
```

If for some reason you want to stop using an `mplib` instance while processing is not yet actually done, you can call `mp:finish`. Eventually, used memory will be freed and open files will be closed by the Lua garbage collector, but an explicit `mp:finish` is the only way to capture the final part of the output streams.



7.9.5 Result table

The return value of `mp:execute` and `mp:finish` is a table with a few possible keys (only `status` is always guaranteed to be present).

<code>log</code>	string	output to the 'log' stream
<code>term</code>	string	output to the 'term' stream
<code>error</code>	string	output to the 'error' stream (only used for 'out of memory')
<code>status</code>	number	the return value: <code>0</code> = good, <code>1</code> = warning, <code>2</code> = errors, <code>3</code> = fatal error
<code>fig</code>	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:

<code>boundingbox</code>	function	returns the bounding box, as an array of 4 values
<code>postscript</code>	function	returns a string that is the ps output of the <code>fig</code> . this function accepts two optional integer arguments for specifying the values of <code>prologues</code> (first argument) and <code>procset</code> (second argument)
<code>svg</code>	function	returns a string that is the svg output of the <code>fig</code> . This function accepts an optional integer argument for specifying the value of <code>prologues</code>
<code>objects</code>	function	returns the actual array of graphic objects in this <code>fig</code>
<code>copy_objects</code>	function	returns a deep copy of the array of graphic objects in this <code>fig</code>
<code>filename</code>	function	the filename this <code>fig</code> 's PostScript output would have written to in stand alone mode
<code>width</code>	function	the <code>fontcharwd</code> value
<code>height</code>	function	the <code>fontcharht</code> value
<code>depth</code>	function	the <code>fontchardp</code> value
<code>italcorr</code>	function	the <code>fontcharit</code> value
<code>charcode</code>	function	the (rounded) <code>charcode</code> 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 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 explicitly 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.



7.9.5.1 fill

path	table	the list of knots
htap	table	the list of knots for the reversed trajectory
pen	table	knots of the pen
color	table	the object's color
linejoin	number	line join style (bare number)
miterlimit	number	miterlimit
prescript	string	the prescript text
postscript	string	the postscript text

The entries `htap` and `pen` are optional.

There is helper function (`mplib.pen_info(obj)`) that returns a table containing a bunch of vital characteristics of the used pen (all values are floats):

width	number	width of the pen
sx	number	x scale
rx	number	xy multiplier
ry	number	yx multiplier
sy	number	y scale
tx	number	x offset
ty	number	y offset

7.9.5.2 outline

path	table	the list of knots
pen	table	knots of the pen
color	table	the object's color
linejoin	number	line join style (bare number)
miterlimit	number	miterlimit
linecap	number	line cap style (bare number)
dash	table	representation of a dash list
prescript	string	the prescript text
postscript	string	the postscript text

The entry `dash` is optional.

7.9.5.3 text

text	string	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

7.9.5.4 special

prescript	string	special text
-----------	--------	--------------

7.9.5.5 start_bounds, start_clip

path	table	the list of knots
------	-------	-------------------

7.9.5.6 stop_bounds, stop_clip

Here are no fields available.

7.9.6 Subsidiary table formats

7.9.6.1 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.

left_type	string	when present: endpoint, but usually absent
right_type	string	like <code>left_type</code>
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.

7.9.6.2 Colors

A color is an integer array with 0, 1, 3 or 4 values:

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`.



7.9.6.3 Transforms

Each transform is a six-item array.

- 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.

7.9.6.4 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.

`dashes` hash an array of on-off numbers
`offset` number the starting offset value

7.9.7 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.

7.9.7.1 `mp:char_width`

```
<number> w = mp:char_width(<string> fontname, <number> char)
```

7.9.7.2 `mp:char_height`

```
<number> w = mp:char_height(<string> fontname, <number> char)
```

7.9.7.3 `mp:char_depth`

```
<number> w = mp:char_depth(<string> fontname, <number> char)
```

7.10 The node library

The `node` library contains functions that facilitate dealing with (lists of) nodes and their values. They allow you to create, alter, copy, delete, and insert Lua_T_E_X node objects, the core objects within the typesetter.

Lua_T_E_X nodes are represented in Lua as userdata with the metadata type `luatex.node`. 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 `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`, but it is most important when dealing with 'whatsits', because they are differentiated solely based on their `subtype`.

The other available fields depend on the `id` (and for 'whatsits', the `subtype`) of the node. Further details on the various fields and their meanings are given in chapter 6.

Support for `unset` (alignment) nodes is partial: they can be queried and modified from Lua code, but not created.

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.

At the moment, memory management of nodes should still 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 LuaTeX 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.

There are statistics available with regards to the allocated node memory, which can be handy for tracing.

7.10.1 Node handling functions

7.10.1.1 `node.is_node`

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

This function returns true if the argument is a userdata object of type `<node>`.

7.10.1.2 `node.types`

```
<table> t = node.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.

7.10.1.3 `node.whatsits`

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



TeX's 'whatsits' all have the same `id`. The various subtypes are defined by their `subtype` fields. The function is much like `node.types`, except that it provides an array of `subtype` mappings.

7.10.1.4 `node.id`

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

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

7.10.1.5 `node.subtype`

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

This converts a single whatsit name to its internal numeric representation (`subtype`).

7.10.1.6 `node.type`

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

In the argument is a number, then this 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.

7.10.1.7 `node.fields`

```
<table> t = node.fields(<number> id)
<table> t = node.fields(<number> id, <number> subtype)
```

This function returns an array of valid field names for a particular type of node. If you want to get the valid fields for a 'whatsit', you have to supply the second argument also. In other cases, any given second argument will be silently ignored.

This function accepts string `id` and `subtype` values as well.

7.10.1.8 `node.has_field`

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

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

7.10.1.9 `node.new`

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

Creates a new node. All of the new node's fields are initialized to either zero or `nil` except for `id` and `subtype` (if supplied). If you want to create a new whatsit, then the second argument is



required, otherwise it need not be present. As with all node functions, this function creates a node on the $\text{T}_{\text{E}}\text{X}$ level.

This function accepts string `id` and `subtype` values as well.

7.10.1.10 `node.free`

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

Removes the node `n` from $\text{T}_{\text{E}}\text{X}$'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.

7.10.1.11 `node.flush_list`

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

Removes the node list `n` and the complete node list following `n` from $\text{T}_{\text{E}}\text{X}$'s memory. 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.

7.10.1.12 `node.copy`

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

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.

7.10.1.13 `node.copy_list`

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

Creates a deep copy of the node list that starts at `n`. If `m` is also given, the copy stops just before node `m`.

Note that you cannot copy attribute lists this way, specialized functions for dealing with attribute lists will be provided later but are not there yet. 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.

7.10.1.14 `node.next`

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

Returns the node following this node, or `nil` if there is no such node.



7.10.1.15 `node.prev`

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

Returns the node preceding this node, or `nil` if there is no such node.

7.10.1.16 `node.current_attr`

```
<node> m = node.current_attr()
```

Returns the currently active list of attributes, if there is one.

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`.

Note: this function is somewhat experimental, and it returns the *actual* attribute list, not a copy thereof. Therefore, changing any of the attributes in the list will change these values for all nodes that have the current attribute list assigned to them.

7.10.1.17 `node.hpack`

```
<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)
```

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.

Caveat: at this moment, 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!

7.10.1.18 `node.vpack`

```
<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)
```

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.

The second return value is the badness of the generated box.

See the description of `node.hpack()` for a few memory allocation caveats.

7.10.1.19 `node.dimensions`

```
<number> w, <number> h, <number> d = node.dimensions(<node> n)
<number> w, <number> h, <number> d = node.dimensions(<node> n, <string> dir)
<number> w, <number> h, <number> d = node.dimensions(<node> n, <node> t)
<number> w, <number> h, <number> d = node.dimensions(<node> n, <node> t, <string>
dir)
```

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, <string> dir)
<number> w, <number> h, <number> d =
  node.dimensions(<number> glue_set, <number> glue_sign,
                 <number> glue_order, <node> n, <node> t)
<number> w, <number> h, <number> d =
  node.dimensions(<number> glue_set, <number> glue_sign,
                 <number> glue_order, <node> n, <node> t, <string> dir)
```

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)
)) }

```

7.10.1.20 `node.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`.

7.10.1.21 `node.slide`

```

<node> m = node.slide(<node> n)

```

Returns the last node of the node list that starts at `n`. As a side-effect, it also creates a reverse chain of `prev` pointers between nodes.

7.10.1.22 `node.tail`

```

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

```

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

7.10.1.23 `node.length`

```

<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.

7.10.1.24 `node.count`

```

<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.

7.10.1.25 `node.traverse`

```
<node> t = 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.

7.10.1.26 `node.traverse_id`

```
<node> t = 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

```

7.10.1.27 `node.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 endnode this helper return that node, else it follows the list and return the next math endnote. If no such node is found `nil` is returned.

7.10.1.28 `node.remove`

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

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.

7.10.1.29 `node.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`.

7.10.1.30 `node.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`.



7.10.1.31 `node.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.

7.10.1.32 `node.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).

7.10.1.33 `node.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).

7.10.1.34 `node.unprotect_glyphs`

```
node.unprotect_glyphs(<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.

7.10.1.35 `node.protect_glyphs`

```
node.protect_glyphs(<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.

7.10.1.36 `node.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.



7.10.1.37 `node.write`

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

This is an experimental function that will append a node list to TeX's 'current list'. The node list is not deep-copied! There is no error checking either!

7.10.1.38 `node.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.

7.10.2 Attribute handling

Attributes appear as linked list of userdata objects in the `attr` field of individual nodes. They can be handled individually, but it is much safer and more efficient to use the dedicated functions associated with them.

7.10.2.1 `node.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`.

7.10.2.2 `node.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. *[needs explanation]*

7.10.2.3 `node.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`.



7.11 The pdf library

This contains variables and functions that are related to the pdf backend.

7.11.1 pdf.mapfile, pdf.mapline

```
pdf.mapfile(<string> map file)
pdf.mapline(<string> map line)
```

These two functions can be used to replace primitives `\pdfmapfile` and `\pdfmapline` from pdf_TE_X. They expect a string as only parameter and have no return value.

The also functions replace the former variables `pdf.pdfmapfile` and `pdf.pdfmapline`.

7.11.2 pdf.catalog, pdf.info, pdf.names, pdf.trailer

These variables offer a read-write interface to the corresponding pdf_TE_X token lists. The value types are strings and they are written out to the pdf file directly after the pdf_TE_X token registers.

The preferred interface is now `pdf.setcatalog`, `pdf.setinfo` `pdf.setnames` and `pdf.settrailer` for setting these properties and `pdf.getcatalog`, `pdf.getinfo` `pdf.getnames` and `pdf.gettrailer` for querying them,

The corresponding ‘pdf’ parameter names `pdf.pdfcatalog`, `pdf.pdfinfo`, `pdf.pdfnames`, and `pdf.pdftrailer` are not available.

7.11.3 pdf.<set/get>pageattributes, pdf.<set/get>pageresources, pdf.<set/get>attributes

These variables offer a read-write interface to related token lists. The value types are strings. The variables have no interaction with the corresponding pdf_TE_X token registers `\pdfpageattr`, `\pdfpageresources`, and `\pdfpagesattr`. They are written out to the pdf file directly after the pdf_TE_X token registers.

The preferred interface is now `pdf.setpageattributes`, `pdf.setpagesattributes` and `pdf.setpageresources` for setting these properties and `pdf.getpageattributes`, `pdf.getpagesattributes` and `pdf.getpageresources` for querying them.

7.11.4 pdf.<set/get>xformattributes, pdf.<set/get>xformresources

These variables offer a read-write interface to related token lists. The value types are strings. The variables have no interaction with the corresponding pdf_TE_X token registers `\pdfxformattr` and `\pdfxformresources`. They are written out to the pdf file directly after the pdf_TE_X token registers.

The preferred interface is now `pdf.setxformattributes` and `pdf.setxformresources` for setting these properties and `pdf.getxformattributes` and `pdf.getxformresources` for querying them.



7.11.5 `pdf.setcompresslevel` and `pdf.setobjcompresslevel`

These two functions set the level of compression. The minimum value is 0, the maximum is 9.

7.11.6 `pdf.setdecimaldigits` and `pdf.getdecimaldigits`

These two functions set the accuracy of floats written to the pdf file. You can set any value but the backend will not go below 3 and above 6.

7.11.7 `pdf.setpkresolution` and `pdf.getpkresolution`

This setter takes two arguments: the resolution and an optional zero or one that indicates if this is a fixed one. The getter returns these two values.

7.11.8 `pdf.lastobj`, `pdf.lastlink`, `pdf.lastannot`, and `pdf.retval`

These status variables are similar to the ones traditionally used at the TeX end.

7.11.9 `pdf.setorigin`, `pdf.getorigin`

This one is used to set the horizontal and/or vertical offset (a traditional backend property).

```
pdf.setorigin() -- sets both to 0pt
pdf.setorigin(tex.sp("1in")) -- sets both to 1in
pdf.setorigin(tex.sp("1in"),tex.sp("1in"))
```

The counterpart of this function returns two values.

7.11.10 `pdf.setlinkmargin`, `pdf.getlinkmargin`, `pdf.setdestmargin`, `pdf.getdestmargin`, `pdf.setthreadmargin`, `pdf.getthreadmargin`, `pdf.setxformmargin`, `pdf.getxformmargin`

These functions can be used to set and retrieve the margins that are added to the natural bounding boxes of the respective objects.

7.11.11 `pdf.h`, `pdf.v`

These are the `h` and `v` values that define the current location on the output page, measured from its lower left corner. The values can be queried using scaled points as units.

```
local h = pdf.h
local v = pdf.v
```

7.11.12 `pdf.getpos`, `pdf.gethpos`, `pdf.getvpos`

These are the function variants of `pdf.h` and `pdf.v`. Sometimes using a function is preferred over a key so this saves wrapping. Also, these functions are faster than the key based access,



as `h` and `v` keys are not real variables but looked up using a metatable call. The `getpos` function returns two values, the other return one.

```
local h, v = pdf.getpos()
```

7.11.13 pdf.hasmatrix, pdf.getmatrix

The current matrix transformation is available via the `getmatrix` command, which returns 6 values: `sx`, `rx`, `ry`, `sy`, `tx`, and `ty`. The `hasmatrix` function returns `true` when a matrix is applied.

```
if pdf.hasmatrix() then
    local sx, rx, ry, sy, tx, ty = pdf.getmatrix()
    -- do something useful or not
end
```

7.11.14 pdf.print

A print function to write stuff to the pdf document that can be used from within a `\lualua` argument. This function is not to be used inside `\directlua` unless you know *exactly* what you are doing.

```
pdf.print(<string> s)
pdf.print(<string> type, <string> s)
```

The optional parameter can be used to mimic the behavior of `\pdfliteral`: the `type` is `direct` or `page`.

7.11.15 pdf.immediateobj

This function creates a pdf object and immediately writes it to the pdf file. It is modelled after pdf_{TEX}'s `\immediate \pdfobj` primitives. All function variants return the object number of the newly generated object.

```
<number> n = pdf.immediateobj(<string> objtext)
<number> n = pdf.immediateobj("file", <string> filename)
<number> n = pdf.immediateobj("stream", <string> streamtext, <string> attrtext)
<number> n = pdf.immediateobj("streamfile", <string> filename, <string> attr-
text)
```

The first version puts the `objtext` raw into an object. Only the object wrapper is automatically generated, but any internal structure (like `<< >>` dictionary markers) needs to be provided by the user. The second version with keyword `"file"` as 1st argument puts the contents of the file with name `filename` raw into the object. The third version with keyword `"stream"` creates a stream object and puts the `streamtext` raw into the stream. The stream length is automatically calculated. The optional `attrtext` goes into the dictionary of that object. The fourth version with keyword `"streamfile"` does the same as the 3rd one, it just reads the stream data raw from a file.



An optional first argument can be given to make the function use a previously reserved pdf object.

```
<number> n = pdf.immediateobj(<integer> n, <string> objtext)
<number> n = pdf.immediateobj(<integer> n, "file", <string> filename)
<number> n = pdf.immediateobj(<integer> n, "stream", <string> streamtext, <string>
attrtext)
<number> n = pdf.immediateobj(<integer> n, "streamfile", <string> filename,
<string> attrtext)
```

7.11.16 pdf.obj

This function creates a pdf object, which is written to the pdf file only when referenced, e.g., by `pdf.refobj()`.

All function variants return the object number of the newly generated object, and there are two separate calling modes.

The first mode is modelled after pdf_TE_X's `\pdfobj` primitive.

```
<number> n = pdf.obj(<string> objtext)
<number> n = pdf.obj("file", <string> filename)
<number> n = pdf.obj("stream", <string> streamtext, <string> attrtext)
<number> n = pdf.obj("streamfile", <string> filename, <string> attrtext)
```

An optional first argument can be given to make the function use a previously reserved pdf object.

```
<number> n = pdf.obj(<integer> n, <string> objtext)
<number> n = pdf.obj(<integer> n, "file", <string> filename)
<number> n = pdf.obj(<integer> n, "stream", <string> streamtext, <string> attr-
text)
<number> n = pdf.obj(<integer> n, "streamfile", <string> filename, <string> at-
trtext)
```

The second mode accepts a single argument table with key-value pairs.

```
<number> n = pdf.obj {
  type          = <string>,
  immediate     = <boolean>,
  objnum        = <number>,
  attr          = <string>,
  compresslevel = <number>,
  objcompression = <boolean>,
  file          = <string>,
  string        = <string>
}
```

The `type` field can have the values `raw` and `stream`, this field is required, the others are optional (within constraints).



Note: this mode makes `pdf.obj` look more flexible than it actually is: the constraints from the separate parameter version still apply, so for example you can't have both `string` and `file` at the same time.

7.11.17 `pdf.refobj`

This function, the Lua version of the `\pdfrefobj` primitive, references an object by its object number, so that the object will be written out.

```
pdf.refobj(<integer> n)
```

This function works in both the `\directlua` and `\latelua` environment. Inside `\directlua` a new whatsit node 'pdf_refobj' is created, which will be marked for flushing during page output and the object is then written directly after the page, when also the resources objects are written out. Inside `\latelua` the object will be marked for flushing.

This function has no return values.

7.11.18 `pdf.reserveobj`

This function creates an empty pdf object and returns its number.

```
<number> n = pdf.reserveobj()  
<number> n = pdf.reserveobj("annot")
```

7.11.19 `pdf.registerannot`

This function adds an object number to the `/Annots` array for the current page without doing anything else. This function can only be used from within `\latelua`.

```
pdf.registerannot (<number> objnum)
```

7.11.20 `pdf.newcolorstack`

This function allocates a new color stack and returns its id. The arguments are the same as for the similar backend extension primitive.

```
pdf.newcolorstack("0 g", "page", true) -- page|direct|origin
```

7.12 The `pdfscanner` library

The `pdfscanner` library allows interpretation of PDF content streams and `/ToUnicode` (cmap) streams. You can get those streams from the `epdf` library, as explained in an earlier section. There is only a single top-level function in this library:

```
pdfscanner.scan (<Object> stream, <table> operatortable, <table> info)
```



The first argument, `stream`, should be either a PDF stream object, or a PDF array of PDF stream objects (those options comprise the possible return values of `<Page>:getContents()` and `<Object>:getStream()` in the `epdf` library).

The second argument, `operatortable`, should be a Lua table where the keys are PDF operator name strings and the values are Lua functions (defined by you) that are used to process those operators. The functions are called whenever the scanner finds one of these PDF operators in the content stream(s). The functions are called with two arguments: the `scanner` object itself, and the `info` table that was passed are the third argument to `pdfscanner.scan`.

Internally, `pdfscanner.scan` loops over the PDF operators in the stream(s), collecting operands on an internal stack until it finds a PDF operator. If that PDF operator's name exists in `operatortable`, then the associated function is executed. After the function has run (or when there is no function to execute) the internal operand stack is cleared in preparation for the next operator, and processing continues.

The `scanner` argument to the processing functions is needed because it offers various methods to get the actual operands from the internal operand stack.

A simple example of processing a PDF's document stream could look like this:

```
function Do (scanner, info)
  local val      = scanner:pop()
  local name     = val[2] -- val[1] == 'name'
  local resources = info.resources
  local xobject  = resources:lookup("XObject"):getDict():lookup(name)
  print (info.space .. 'Use XObject ' .. name)
  if xobject and xobject:isStream() then
    local dict = xobject:getStream():getDict()
    if dict then
      local name = dict:lookup("Subtype")
      if name:getName() == "Form" then
        local newinfo = {
          space = info.space .. " " ,
          resources = dict:lookup("Resources"):getDict()
        }
        pdfscanner.scan(xobject, operatortable, newinfo)
      end
    end
  end
end

operatortable = { Do = Do }

doc      = epdf.open(arg[1])
pagenum = 1

while pagenum <= doc:getNumPages() do
  local page = doc:getCatalog():getPage(pagenum)
```



```

local info = {
    space      = "  ",
    resources  = page:getResourceDict()
}
print('Page ' .. pagenum)
pdfscanner.scan(page:getContents(), operatortable, info)
pagenum = pagenum + 1
end

```

This example iterates over all the actual content in the PDF, and prints out the found XObject names. While the code demonstrates quite some of the `epdf` functions, let's focus on the type `pdfscanner` specific code instead.

From the bottom up, the line

```
pdfscanner.scan(page:getContents(), operatortable, info)
```

runs the scanner with the PDF page's top-level content.

The third argument, `info`, contains two entries: `space` is used to indent the printed output, and `resources` is needed so that embedded XForms can find their own content.

The second argument, `operatortable` defines a processing function for a single PDF operator, `Do`.

The function `Do` prints the name of the current XObject, and then starts a new scanner for that object's content stream, under the condition that the XObject is in fact a `/Form`. That nested scanner is called with new `info` argument with an updated `space` value so that the indentation of the output nicely nests, and with a new `resources` field to help the next iteration down to properly process any other, embedded XObjects.

Of course, this is not a very useful example in practise, but for the purpose of demonstrating `pdfscanner`, it is just long enough. It makes use of only one `scanner` method: `scanner:pop()`. That function pops the top operand of the internal stack, and returns a Lua table where the object at index one is a string representing the type of the operand, and object two is its value.

The list of possible operand types and associated Lua value types is:

```

integer  <number>
real     <number>
boolean  <boolean>
name     <string>
operator <string>
string   <string>
array    <table>
dict     <table>

```

In case of `integer` or `real`, the value is always a Lua (floating point) number.

In case of `name`, the leading slash is always stripped.

In case of `string`, please bear in mind that PDF actually supports different types of strings (with different encodings) in different parts of the PDF document, so may need to reencode some of the



results; `pdfscanner` always outputs the byte stream without reencoding anything. `pdfscanner` does not differentiate between literal strings and hexadecimal strings (the hexadecimal values are decoded), and it treats the stream data for inline images as a string that is the single operand for `EI`.

In case of `array`, the table content is a list of `pop` return values.

In case of `dict`, the table keys are PDF name strings and the values are `pop` return values.

There are few more methods defined that you can ask `scanner`:

<code>pop</code>	as explained above
<code>popNumber</code>	return only the value of a <code>real</code> or <code>integer</code>
<code>popName</code>	return only the value of a <code>name</code>
<code>popString</code>	return only the value of a <code>string</code>
<code>popArray</code>	return only the value of a <code>array</code>
<code>popDict</code>	return only the value of a <code>dict</code>
<code>popBool</code>	return only the value of a <code>boolean</code>
<code>done</code>	abort further processing of this <code>scan()</code> call

The `popXXX` are convenience functions, and come in handy when you know the type of the operands beforehand (which you usually do, in PDF). For example, the `Do` function could have used `local name = scanner:popName()` instead, because the single operand to the `Do` operator is always a PDF name object.

The `done` function allows you to abort processing of a stream once you have learned everything you want to learn. This comes in handy while parsing `/ToUnicode`, because there usually is trailing garbage that you are not interested in. Without `done`, processing only end at the end of the stream, possibly wasting CPU cycles.

7.13 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>pdf_gone</code>	written pdf bytes
<code>pdf_ptr</code>	not yet written pdf bytes
<code>dvi_gone</code>	written dvi bytes
<code>dvi_ptr</code>	not yet written dvi bytes
<code>total_pages</code>	number of written pages



output_file_name	name of the pdf or dvi file
log_name	name of the log file
banner	terminal display banner
var_used	variable (one-word) memory in use
dyn_used	token (multi-word) memory in use
str_ptr	number of strings
init_str_ptr	number of ini \TeX strings
max_strings	maximum allowed strings
pool_ptr	string pool index
init_pool_ptr	ini \TeX string pool index
pool_size	current size allocated for string characters
node_mem_usage	a string giving insight into currently used nodes
var_mem_max	number of allocated words for nodes
fix_mem_max	number of allocated words for tokens
fix_mem_end	maximum number of used tokens
cs_count	number of control sequences
hash_size	size of hash
hash_extra	extra allowed hash
font_ptr	number of active fonts
input_ptr	th elevel of input we're at
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_buf_stack	max used buffer position
max_save_stack	max used save stack entries
stack_size	input stack size
nest_size	nesting stack size
param_size	parameter stack size
buf_size	current allocated size of the line buffer
save_size	save stack size
obj_ptr	max pdf object pointer
obj_tab_size	pdf object table size
pdf_os_cntr	max pdf object stream pointer
pdf_os_objidx	pdf object stream index
pdf_dest_names_ptr	max pdf destination pointer
dest_names_size	pdf destination table size
pdf_mem_ptr	max pdf memory used
pdf_mem_size	pdf memory size
largest_used_mark	max referenced marks class
filename	name of the current input file
inputid	numeric id of the current input
linenumber	location in the current input file
lasterrorstring	last \TeX error string
lastluaerrorstring	last Lua error string
lastwarningtag	last warning string
lastwarningstring	last warning tag, normally an indication of in what part



<code>lasterrorcontext</code>	last error context string (with newlines)
<code>luabytecodes</code>	number of active Lua bytecode registers
<code>luabytecode_bytes</code>	number of bytes in Lua bytecode registers
<code>luastate_bytes</code>	number of bytes in use by Lua interpreters
<code>output_active</code>	<code>true</code> if the <code>\output</code> routine is active
<code>callbacks</code>	total number of executed callbacks so far
<code>indirect_callbacks</code>	number of those that were themselves a result of other callbacks (e.g. file readers)
<code>luatex_version</code>	the LuaTeX version number
<code>luatex_revision</code>	the LuaTeX revision string
<code>ini_version</code>	<code>true</code> if this is an iniTeX run
<code>shell_escape</code>	<code>0</code> means disabled, <code>1</code> is restricted and <code>2</code> means anything is permitted

The error and warning messages can be wiped with the `resetmessages` function.

7.14 The `tex` library

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 has these characteristics:

- You can use it after `\the`
- It is a single token.
- Some special others, see the list below

This excludes parameters that need extra arguments, like `\the\scriptfont`.

The subset comprising simple integer and dimension registers are writable as well as readable (stuff like `\tracingcommands` and `\parindent`).

7.14.1 Internal parameter values

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`. If you created aliases, you can use accessors like `tex.getdimen` as these also understand names of built-in variables.

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 (<string> n, ...)
tex.set ("global", <string> n, ...)
... = tex.get (<string> n)
```

There are also dedicated setters, getters and checkers:



```

local d = tex.getdimen("foo")
if tex.isdimen("bar") then
    tex.setdimen("bar",d)
end

```

There are such helpers for `dimen`, `count`, `skip`, `box` and `attribute` registers.

7.14.1.1 Integer parameters

The integer parameters accept and return Lua numbers.

Read-write:

<code>tex.adjdemerits</code>	<code>tex.newlinechar</code>
<code>tex.binoppenalty</code>	<code>tex.outputpenalty</code>
<code>tex.brokenpenalty</code>	<code>tex.pausing</code>
<code>tex.catcodetable</code>	<code>tex.postdisplaypenalty</code>
<code>tex.clubpenalty</code>	<code>tex.predisplaydirection</code>
<code>tex.day</code>	<code>tex.predisplaypenalty</code>
<code>tex.defaultthyphenchar</code>	<code>tex.pretolerance</code>
<code>tex.defaultskewchar</code>	<code>tex.relpnalty</code>
<code>tex.delimiterfactor</code>	<code>tex.righthyphenmin</code>
<code>tex.displaywidowpenalty</code>	<code>tex.savinghyphcodes</code>
<code>tex.doublehyphdemerits</code>	<code>tex.savingvdiscards</code>
<code>tex.endlinechar</code>	<code>tex.showboxbreadth</code>
<code>tex.errorcontextlines</code>	<code>tex.showboxdepth</code>
<code>tex.escapechar</code>	<code>tex.time</code>
<code>tex.exhyphenpenalty</code>	<code>tex.tolerance</code>
<code>tex.fam</code>	<code>tex.tracingassigns</code>
<code>tex.finalhyphdemerits</code>	<code>tex.tracingcommands</code>
<code>tex.floatingpenalty</code>	<code>tex.tracinggroups</code>
<code>tex.globaldefs</code>	<code>tex.tracingifs</code>
<code>tex.hangafter</code>	<code>tex.tracinglostchars</code>
<code>tex.hbadness</code>	<code>tex.tracingmacros</code>
<code>tex.holdinginserts</code>	<code>tex.tracingnesting</code>
<code>tex.hyphenpenalty</code>	<code>tex.tracingonline</code>
<code>tex.interlinepenalty</code>	<code>tex.tracingoutput</code>
<code>tex.language</code>	<code>tex.tracingpages</code>
<code>tex.lastlinefit</code>	<code>tex.tracingparagraphs</code>
<code>tex.lefthyphenmin</code>	<code>tex.tracingrestores</code>
<code>tex.linepenalty</code>	<code>tex.tracingscantokens</code>
<code>tex.localbrokenpenalty</code>	<code>tex.tracingstats</code>
<code>tex.localinterlinepenalty</code>	<code>tex.uchyph</code>
<code>tex.looseness</code>	<code>tex.vbadness</code>
<code>tex.mag</code>	<code>tex.widowpenalty</code>
<code>tex.maxdeadcycles</code>	<code>tex.year</code>
<code>tex.month</code>	



Read-only:

<code>tex.deadcycles</code>	<code>tex.parshape</code>	<code>tex.spacefactor</code>
<code>tex.insertpenalties</code>	<code>tex.prevgraf</code>	

7.14.1.2 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.

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>

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 \TeX 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 Lua \TeX can guess



your intentions. By using the accessor in the `tex` tables, you get and set the values at the top of the nest stack.

7.14.1.3 Direction parameters

The direction parameters are read-only and return a Lua string.

<code>tex.bodydir</code>	<code>tex.pagedir</code>	<code>tex.textdir</code>
<code>tex.mathdir</code>	<code>tex.pardir</code>	

7.14.1.4 Glue parameters

The glue parameters accept and return a userdata object that represents a `glue_spec` node.

<code>tex.abovedisplayshortskip</code>	<code>tex.leftskip</code>	<code>tex.spaceskip</code>
<code>tex.abovedisplayskip</code>	<code>tex.lineskip</code>	<code>tex.splittopskip</code>
<code>tex.baselineskip</code>	<code>tex.parfillskip</code>	<code>tex.tabskip</code>
<code>tex.belowdisplayshortskip</code>	<code>tex.parskip</code>	<code>tex.topskip</code>
<code>tex.belowdisplayskip</code>	<code>tex.rightskip</code>	<code>tex.xspaceskip</code>

7.14.1.5 Muglue parameters

All muglue parameters are to be used read-only and return a Lua string.

<code>tex.medmuskip</code>	<code>tex.thickmuskip</code>	<code>tex.thinmuskip</code>
----------------------------	------------------------------	-----------------------------

7.14.1.6 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.

<code>tex.errhelp</code>	<code>tex.everyjob</code>	<code>tex.pdfpageattr</code>
<code>tex.everycr</code>	<code>tex.everymath</code>	<code>tex.pdfpageresources</code>
<code>tex.everydisplay</code>	<code>tex.everypar</code>	<code>tex.pdfpagesattr</code>
<code>tex.everyeof</code>	<code>tex.everyvbox</code>	<code>tex.pdfpkmode</code>
<code>tex.everyhbox</code>	<code>tex.output</code>	

7.14.2 Convert commands

All 'convert' commands are read-only and return a Lua string. The supported commands at this moment are:

<code>tex.eTeXVersion</code>	<code>tex.formatname</code>
<code>tex.eTeXrevision</code>	<code>tex.jobname</code>



<code>tex.luatexbanner</code>	<code>tex.uniformdeviate(number)</code>
<code>tex.luatexrevision</code>	<code>tex.number(number)</code>
<code>tex.pdfnormaldeviate</code>	<code>tex.romannumeral(number)</code>
<code>tex.fontname(number)</code>	<code>tex.pdfpageref(number)</code>
<code>tex.pdffontname(number)</code>	<code>tex.pdfxformname(number)</code>
<code>tex.pdffontobjnum(number)</code>	<code>tex.fontidentifier(number)</code>
<code>tex.pdffontsize(number)</code>	

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).

if these are really needed in a macro package.

7.14.3 Last item commands

All ‘last item’ commands are read-only and return a number.

The supported commands at this moment are:

<code>tex.lastpenalty</code>	<code>tex.pdflastximagepages</code>	<code>tex.eTeXversion</code>
<code>tex.lastkern</code>	<code>tex.pdflastannot</code>	<code>tex.currentgrouplevel</code>
<code>tex.lastskip</code>	<code>tex.pdflastxpos</code>	<code>tex.currentgrouptype</code>
<code>tex.lastnodetype</code>	<code>tex.pdflastypos</code>	<code>tex.currentiflevel</code>
<code>tex.inputlineno</code>	<code>tex.pdfrandomseed</code>	<code>tex.currentiftype</code>
<code>tex.pdflastobj</code>	<code>tex.pdflastlink</code>	<code>tex.currentifbranch</code>
<code>tex.pdflastxform</code>	<code>tex.luatexversion</code>	<code>tex.pdflastximagecol-</code>
<code>tex.pdflastximage</code>	<code>tex.eTeXminorversion</code>	<code>ordepth</code>

7.14.4 Attribute, count, dimension, skip and token registers

TeX’s attributes (`\attribute`), counters (`\count`), dimensions (`\dimen`), skips (`\skip`) and token (`\toks`) registers can be accessed and written to using two times five virtual sub-tables of the `tex` table:

<code>tex.attribute</code>	<code>tex.dimen</code>	<code>tex.toks</code>
<code>tex.count</code>	<code>tex.skip</code>	

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, LuaTeX 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 attribute and 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).

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 (<number> n, <node> s)
tex.setskip (<string> s, <node> s)
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 muskips by nodes, and toks by strings. For tokens registers we have an alternative where a catcode table is specified:

```
tex.scantoks(0,3,"$e=mc^2$")
tex.scantoks("global",0,"$\int\limits^1_2$")
```

In the function-based interface, it is possible to define values globally by using the string `global` as the first function argument.

7.14.5 Character code registers

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

```
tex.lccode          tex.sfcode          tex.mathcode
tex.uccode          tex.catcode          tex.delcode
```

The function call interfaces are roughly as above, but there are a few twists. `sfcodes` are the simple ones:

```
tex.setsfcode (<number> n, <number> s)
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> mathclass, <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:

```
class, family, char = tex.getmathcodes (<number> n)
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`.

7.14.6 Box registers

It is possible to set and query actual boxes, using the node interface as defined in the `node` library:

```
tex.box
```

for array access, or



```

tex.setbox(<number> n, <node> s)
tex.setbox(<string> cs, <node> s)
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 T_EX 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])
```

The following function will register a box for reuse (this is modelled after so called xforms in pdf). You can (re)use the box with `\useboxresource` or by creating a rule node with subtype 2.

```
local index = tex.saveboxresource(n, attributes, resources, immediate)
```

The optional second and third arguments are strings, the fourth is a boolean.

You can generate the reference (a rule type) with:

```
local reused = tex.useboxresource(n, wd, ht, dp)
```

The dimensions are optional and the final ones are returned as extra values. The following is just a bonus (no dimensions returned means that the resource is unknown):

```
local w, h, d = tex.getboxresourcedimensions(n)
```

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.

7.14.7 Math parameters

It is possible to set and query the internal math parameters using:

```

tex.setmath(<string> n, <string> t, <number> n)
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:

<code>quad</code>	<code>axis</code>	<code>operatorsize</code>	
<code>overbarkern</code>	<code>overbarrule</code>	<code>overbarvgap</code>	
<code>underbarkern</code>	<code>underbarrule</code>	<code>underbarvgap</code>	
<code>radicalkern</code>	<code>radicalrule</code>	<code>radicalvgap</code>	
<code>radicaldegreebefore</code>	<code>radicaldegreeafter</code>	<code>radicaldegreeraise</code>	
<code>stackvgap</code>	<code>stacknumup</code>	<code>stackdenomdown</code>	
<code>fractionrule</code>	<code>fractionnumvgap</code>	<code>fractionnumup</code>	
<code>fractiondenomvgap</code>	<code>fractiondenomdown</code>	<code>fractiondelsize</code>	
<code>limitabovevgap</code>	<code>limitabovebgap</code>	<code>limitabovekern</code>	
<code>limitbelowvgap</code>	<code>limitbelowbgap</code>	<code>limitbelowkern</code>	
<code>underdelimitervgap</code>	<code>underdelimiterbgap</code>		
<code>overdelimitervgap</code>	<code>overdelimiterbgap</code>		
<code>subshiftdrop</code>	<code>supshiftdrop</code>	<code>subshiftdown</code>	
<code>subsupshiftdown</code>	<code>subtopmax</code>	<code>supshiftdown</code>	
<code>supbottommin</code>	<code>supsubbottommax</code>	<code>subsupvgap</code>	
<code>spaceafterscript</code>	<code>connectoroverlapmin</code>		
<code>ordordspacing</code>	<code>ordopspacing</code>	<code>ordbinspacing</code>	<code>ordrelspacing</code>
<code>ordopenspacing</code>	<code>ordclosespacing</code>	<code>ordpunctspacing</code>	<code>ordinnerspacing</code>
<code>opordspacing</code>	<code>opopspacing</code>	<code>opbinspacing</code>	<code>oprelspacing</code>
<code>opopenspacing</code>	<code>opclosespacing</code>	<code>oppunctspacing</code>	<code>opinnerspacing</code>
<code>binordspacing</code>	<code>binopspacing</code>	<code>binbinspacing</code>	<code>binrelspacing</code>
<code>binopenspacing</code>	<code>binclosespacing</code>	<code>binpunctspacing</code>	<code>bininnerspacing</code>
<code>relordspacing</code>	<code>relopspacing</code>	<code>relbinspacing</code>	<code>relrelspacing</code>
<code>relopenspacing</code>	<code>relclosespacing</code>	<code>relpunctspacing</code>	<code>relinnerspacing</code>
<code>openordspacing</code>	<code>openopspacing</code>	<code>openbinspacing</code>	<code>openrelspacing</code>
<code>openopenspacing</code>	<code>openclosespacing</code>	<code>openpunctspacing</code>	<code>openinnerspacing</code>
<code>closeordspacing</code>	<code>closeopspacing</code>	<code>closebinspacing</code>	<code>closerelspacing</code>
<code>closeopenspacing</code>	<code>closeclosespacing</code>	<code>closepunctspacing</code>	<code>closeinnerspacing</code>
<code>punctordspacing</code>	<code>punctopspacing</code>	<code>punctbinspacing</code>	<code>punctrelspacing</code>
<code>punctopenspacing</code>	<code>punctclosespacing</code>	<code>punctpunctspacing</code>	<code>punctinnerspacing</code>
<code>innerordspacing</code>	<code>inneropspacing</code>	<code>innerbinspacing</code>	<code>innerrelspacing</code>
<code>inneropenspacing</code>	<code>innerclosespacing</code>	<code>innerpunctspacing</code>	<code>innerinnerspacing</code>

The values for the style parameter name are:

<code>display</code>	<code>crampeddisplay</code>
<code>text</code>	<code>crampedtext</code>
<code>script</code>	<code>crampedscript</code>
<code>scriptscript</code>	<code>crampedscriptscript</code>

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.



7.14.8 Special list heads

The virtual table `tex.lists` contains the set of internal registers that keep track of building page lists.

field	description
<code>page_ins_head</code>	circular list of pending insertions
<code>contrib_head</code>	the recent contributions
<code>page_head</code>	the current page content
<code>hold_head</code>	used for held-over items for next page
<code>adjust_head</code>	head of the current <code>\vadjust</code> list
<code>pre_adjust_head</code>	head of the current <code>\vadjust pre</code> list
<code>page_discards_head</code>	head of the discarded items of a page break
<code>split_discards_head</code>	head of the discarded items in a vsplit

7.14.9 Semantic nest levels

The virtual table `tex.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 `tex.nest.ptr`, which gives the highest available index. Neither the array items in `tex.nest[]` nor `tex.nest.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, `tex.nest[0]` the outermost (main vertical list) level.

The known fields are:

key	type	modes	explanation
<code>mode</code>	number	all	The current mode. This is a number representing the main mode at this level: <code>0</code> == no mode (this happens during <code>\write</code>) <code>1</code> == vertical, <code>127</code> = horizontal, <code>253</code> = display math. <code>-1</code> == internal vertical, <code>-127</code> = restricted horizontal, <code>-253</code> = inline math.
<code>modeline</code>	number	all	source input line where this mode was entered in, negative inside the output routine
<code>head</code>	node	all	the head of the current list
<code>tail</code>	node	all	the tail of the current list
<code>prevgraf</code>	number	vmode	number of lines in the previous paragraph
<code>prevdepth</code>	number	vmode	depth of the previous paragraph (equal to <code>\pdfignoreddimen</code> when it is to be ignored)
<code>spacefactor</code>	number	hmode	the current space factor
<code>dirs</code>	node	hmode	used for temporary storage by the line break algorithm



<code>noad</code>	node	mmode	used for temporary storage of a pending fraction numerator, for <code>\over</code> etc.
<code>delimptr</code>	node	mmode	used for temporary storage of the previous math delimiter, for <code>\middle</code>
<code>mathdir</code>	boolean	mmode	true when during math processing the <code>\mathdir</code> is not the same as the surrounding <code>\textdir</code>
<code>mathstyle</code>	number	mmode	the current <code>\mathstyle</code>

7.14.10 Print functions

The `tex` table also contains the three print functions that are the major interface from Lua scripting to T_EX.

The arguments to these three functions are all stored in an in-memory virtual file that is fed to the T_EX 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 T_EX's input buffer.

The result of using these functions from inside callbacks is undefined at the moment.

7.14.10.1 `tex.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 T_EX 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 `\endlinechar` appended, all others do.

7.14.10.2 `tex.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 T_EX as a special kind of input line that makes it suitable for use as a partial line input mechanism:



- T_EX does not switch to the ‘new line’ state, so that leading spaces are not ignored.
 - No `\endlinechar` is inserted.
 - Trailing spaces are not removed.
- Note that this does not prevent T_EX 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()`.

7.14.10.3 `tex.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.

7.14.10.4 `tex.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
```

7.14.10.5 `tex.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).



7.14.11 Helper functions

7.14.11.1 `tex.round`

```
<number> n = tex.round(<number> o)
```

Rounds Lua number `o`, and returns a number that is in the range of a valid \TeX register value. If the number starts out of range, it generates a ‘number to big’ error as well.

7.14.11.2 `tex.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 \TeX 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 to 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 Lua \TeX ’s internal, 100% portable scale function will be added at a later date.

7.14.11.3 `tex.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 Lua \TeX would use if it was scanning a dimension specifier in its \TeX -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)

7.14.11.4 `tex.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).



7.14.11.5 `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.

7.14.11.6 `tex.error`

```
tex.error(<string> s)
tex.error(<string> s, <table> help)
```

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.

7.14.11.7 `tex.hashtokens`

```
for i,v in pairs (tex.hashtokens()) do ... end
```

Returns a name and token table pair (see section ?? about token tables) iterator for every non-zero entry in the hash table. 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.

7.14.12 Functions for dealing with primitives

7.14.12.1 `tex.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 T_EX82, 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 Lua_{TEX} is run with `--ini` only the \TeX 82 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.extraprimatives ())}
\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 \TeX 82-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.

7.14.12.2 `tex.extraprimatives`

```
<table> t = tex.extraprimatives(<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:

name values

```
tex      vskip write vsize \normalcontrolspace boundary unhcopy output - / unskip
         unvbox boxmaxdepth muskipdef string toksdef floatingpenalty righthyphen-
         min voffset escapechar topmark splitfirstmark vsplit everydisplay badness
         xleaders textfont showlists language mathchoice topskip abovedisplayshort-
         skip underline tracinglostchars pagefillstretch unvcopy splitbotmark fi-
         nalhyphendemerits atopwithdelims pretolerance fi dp setlanguage ht math-
         chardef nulldelimiterspace or wd pagegoal advance chardef catcode mathchar
         scriptscriptfont mathcode leftskip pageshrink pagefilstretch delcode font-
         name brokenpenalty lastkern belowdisplayshortskip tolerance mathopen ex-
         hyphenpenalty maxdepth futurelet abovewithdelims csstring hangindent last-
         skip linepenalty everyjob xspaceskip globaldefs everypar scriptfont delim-
         iter afterassignment firstmark lineskiplimit lineskip def fam day iffalse
         textstyle end mag box belowdisplayskip ifx let errmessage exhyphenchar hss
```



expandafter the displaywidth Uright mathsurround pagedepth looseness lead-
 ers vss ifhmode botmark ifinner displaystyle accent immediate ifmmode par-
 shape meaning abovedisplayskip medmuskip emergencystretch rightskip math-
 close hangafter hoffset aftergroup cleaders romannumeral hbadness mathbin
 showboxbreadth ifvmode jobname vbadness patterns nonstopmode errhelp predis-
 playpenalty endlinedchar mathinner lastbox showboxdepth postdisplaypenalty
 mathrel holdinginserts radical mathord pagetotal everycr adjdemerits halign
 defaultskewchar errorcontextlines splitmaxdepth Uleft ifcase noindent trac-
 ingmacros moveright predisplaysize tracingrestores message ifhbox deadcy-
 cles interlinepenalty mathpunct lccode noboundary displayindent nonscript
 everyhbox global penalty tracingcommands everymath nolimits noalign inputli-
 neno pagestretch parskip indent dimendef widowpenalty ifvbox above spaceskip
 middle displaylimits pausing everyvbox iftrue moveleft mathop endcsname di-
 men ifcat clubpenalty splittopskip doublehyphendemerits ifdim limits ifeof
 ignorespaces insert delimitershortfall ifodd insertpenalties tracingpages
 hpack vadjust tracingonline count ifnum edef char begingroup sfcodes trac-
 ingparagraphs hyphenation ucode hfuzz openout leqno hyphenpenalty vcenter
 hfil thickmuskip maxdeadcycles mkern hbox overfullrule else hsize raise thin-
 muskip spacefactor input hrule left eqno parfillskip font valign dump relax
 prevdepth read shipout batchmode right skipdef setbox baselineskip special
 mskip endgroup uchyph binoppenalty endinput omit pagefillllstretch overwith-
 delims newlinechar vfilneg time tpack skip vfill span prevgraf over show vbox
 tracingstats year defaultthyphenchar nullfont muskip vpack toks outer multi-
 ply tracingoutput firstvalidlanguage parindent displaywidowpenalty unhbox
 lefthyphenmin vtop mathaccent discretionary vfuzz overline unkern closeout
 showthe showbox uppercase lowercase closein openin errorstopmode scrollmode
 skewchar hyphenchar countdef xdef gdef long Umiddle atop scriptscriptstyle
 scriptstyle unpenalty copy lower kern vfil hfilneg hfill hskip crcr cr ifvoid
 if number lastpenalty par vrule noexpand mark fontdimen divide csname script-
 space outputpenalty month delimiterfactor relpenalty tabskip

core directlua

etex unless botmarks currentiftypet pagediscards mutogluet displaywidowpenalties
 fontcharic fontcharidp fontcharht fontcharwd widowpenalties tracingifs if-
 fontchar eTeXVersion protected topmarks showgroups glueexpr splitfirstmarks
 predisplaydirection everyeof eTeXversion clubpenalties savingvdiscards
 splitbotmarks showtokens tracingassigns dimexpr parshapedimen readline trac-
 ingscantokens tracingnesting ifdefined currentifbranch firstmarks lastnode-
 type marks currentgrouplevel interlinepenalties muexpr unexpanded ifcsname
 parshapeindent showifs parshapelength currentgrouptype splitdiscards glue-
 tomu glueshrink gluestretch glueshrinkorder gluestretchorder numexpr scant-
 okens interactionmode detokenize currentiflevel savinghyphcodes lastlinefit
 tracinggroups eTeXrevision eTeXminorversion

luatex Umathcloseopspacing textdir Umathordpunctspacing Udelimiterunder Uskewed-
 withdelims Umathopenpunctspacing pagebottomoffset mathsurroundskip Umath-
 ordinerspacing Umathbinclosespacing toksapp rightghost Umathlimitbelowgap
 Umathopeninnerspacing tokspre Uoverdelimiter Umathpunctpunctspacing Umath-



closepunctspacing mathdisplayskipmode saveimageresource Umathrelordspacing
 Umathsupbottommin Umathlimitbelowkern copyfont Umathstackdenomdown lo-
 calrightbox Umathfractionrule Umathcharfam Umathcloseinnerspacing Umath-
 openrelspacing Uhexensible Umathsupsubbottommax leftmarginkern Umath-
 closerelspacing ifincsname Umathcharnum Umathinnerordspacing syntex for-
 matname letterspacefont pdfextension Umathrelinnerspacing Umathsubtopmax
 randomseed suppressoutererror Umathsubsupshiftdown Umathopbinspacing Umath-
 ordbinspacing Umathrelopspacing Umathopenbinspacing Umathoverdelimiterbgap
 localleftbox alignmark Uunderdelimiter hyphenationmin Umathclosebinspac-
 ing Umathcodenum dvifedback outputmode luafunction Umathpunctopenspacing
 Umathconnectoroverlapmin crampedscriptscriptstyle Umathradicaldegreearafter
 uniformdeviate luatexversion Umathfractionnumup rightmarginkern Umathop-
 closespacing Umathordclosespacing Umathoverdelimitervgap etokspre expanded
 suppressmathparerror Udelcode bodydir Umathopenclosespacing attribute
 Umathsubshiftdrop Umathsubshiftdown matheqnogapstep Umathpunctrelspacing
 lastsavedimageresourceindex lastsavedimageresourcepages mathoption Umath-
 radicaldegreerise adjustspacing Umathsupshiftdrop Umathcharslot Umath-
 closeclosespacing luatexrevision insertht localinterlinepenalty useboxre-
 source Umathchar Udelimiterover Ustack Umathcode saveboxresource Udelcode-
 num suppresslongerror ignoreligaturesinfont Umathaxis Umathfractionnumvgap
 Umathskewedfractionhgap Umathrelclosespacing Umathpunctbinspacing luatex-
 datestamp Ustopdisplaymath quitvmode crampedscriptstyle letcharcode setran-
 domseed crampedtextstyle pagedir Umathbinrelspacing Umathopordspacing dvi-
 variable attributedef Umathordordspacing pdffeedback Umathskewedfraction-
 vgap Umathopenordspacing mathdir outputbox Umathcloseordspacing pagewidth
 Ustopmath aligntab prehyphenchar dviextension Umathpunctopspacing Umath-
 subsupvgap luaescapestring Umathfractiondenomvgap begincsname Umathradi-
 calrule Umathunderbarrule postexhyphenchar Umathradicaldegreebefore Umath-
 stacknumup normaldeviate Umathbinopspacing boxdir Ustartdisplaymath save-
 catcodetable Umathbinpunctspacing tagcode Uroot lastsavedboxresourcein-
 dex Umathoverbarkern Umathoperatorsize Uradical mathstyle Umathopopenspac-
 ing Umathordopenspacing Umathbininnerspacing Umathinnerrelspacing clear-
 marks Umathoverbarvgap fontid Umathopenopenspacing Umathunderdelimiterbgap
 Umathoverbarrule setfontid crampeddisplaystyle ifabsdim Umathlimitabove-
 bgap Umathcharclass Umathstackvgap Umathinneropspacing Umathrelbinspac-
 ing Umathcloseopenspacing pardir initcatcodetable nokerns pageleftoffset
 tracingfonts nospaces Umathreloppspacing Umathlimitabovekern Udelimiter
 savepos nohrule localbrokenpenalty Umathfractiondelsize gleaders Umath-
 underdelimitervgap Umathinnerbinspacing noligs pagerightoffset draftmode
 Usubscript Umathcharnumdef rrcode Umathaccent pagetopoffset pageheight cat-
 codetable Umathspaceafterscript primitive Umathinneropspacing Uskewed
 pxdimen Umathordopspacing Umathopenopspacing ifabsnum scantexttokens math-
 scriptsmode suppressifcsnameerror suppressfontnotfounderror pdfvariable
 latelua useimageresource efcodes lpcodes hjcodes preexhyphenchar posthyphen-
 char Umathinnerinnerspacing Umathinnerpunctspacing Umathinnerclosespacing
 Umathpunctinnerspacing Umathpunctclosespacing Umathpunctordspacing Umath-



relpunctspacing Umathrelrelspacing Umathbinopenspacing Umathbinbinspacing
 Umathbinordspacing Umathopinnerspacing Umathoppunctspacing Umathoprelspac-
 ing Umathopopspacing Umathordrelspacing Umathsupshiftup Umathlimitbelowvgap
 Umathlimitabovevgap Umathfractiondenomdown Umathradicalvgap Umathradicalk-
 ern Umathunderbarvgap Umathunderbarkern Umathquad Umathchardef Uvextensi-
 ble U superscript Ustartmath ifprimitive Uchar luatexbanner lastypos lastxpos
 novrule etoksapp leftghost expandglyphsinfont lastnamedcs protrudechars

Note that 'luatex' 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' and 'luatex')`.

7.14.12.3 tex.primitives

```
<table> t = tex.primitives()
```

This function returns a hash table listing all primitives that Lua \TeX knows about. The keys in the hash are primitives names, the values are tables representing tokens (see section ??). The third value is always zero.

In the beginning we had `omega` and `pdftex` subsets but in the meantime relevant primitives ave been promoted (either or not adapted) to the `luatex` set when found useful, or removed when considered to be of no use. Originally we had two sets of math definition primitives but the Omega ones have been removed, so we no longer have a subset for math either.

7.14.13 Core functionality interfaces

7.14.13.1 tex.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 $100(t/s)^3$;

7.14.13.2 tex.resetparagraph

This function resets the parameters that \TeX normally resets when a new paragraph is seen.

7.14.13.3 tex.linebreak

```
local <node> nodelist, <table> info =  
    tex.linebreak(<node> listhead, <table> parameters)
```

The understood parameters are as follows:



name	type	description
<code>pardir</code>	string	
<code>pretolerance</code>	number	
<code>tracingparagraphs</code>	number	
<code>tolerance</code>	number	
<code>looseness</code>	number	
<code>hyphenpenalty</code>	number	
<code>exhyphenpenalty</code>	number	
<code>pdfadjustspacing</code>	number	
<code>adjdemerits</code>	number	
<code>pdfprotrudechars</code>	number	
<code>linepenalty</code>	number	
<code>lastlinefit</code>	number	
<code>doublehyphendemerits</code>	number	
<code>finalhyphendemerits</code>	number	
<code>hangafter</code>	number	
<code>interlinepenalty</code>	number or table	if a table, then it is an array like <code>\interlinepenalties</code>
<code>clubpenalty</code>	number or table	if a table, then it is an array like <code>\clubpenalties</code>
<code>widowpenalty</code>	number or table	if a table, then it is an array like <code>\widowpenalties</code>
<code>brokenpenalty</code>	number	
<code>emergencystretch</code>	number	in scaled points
<code>hangindent</code>	number	in scaled points
<code>hsize</code>	number	in scaled points
<code>leftskip</code>	glue_spec node	
<code>rightskip</code>	glue_spec node	
<code>pdfignoreddimen</code>	number	in scaled points
<code>parshape</code>	table	

Note that there is no interface for `\displaywidowpenalties`, you have to pass the right choice for `widowpenalties` yourself.

The meaning of the various keys should be fairly obvious from the table (the names match the \TeX and pdf \TeX primitives) except for the last 5 entries. The four `pdf...line...` keys are ignored if their value equals `pdfignoreddimen`.

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:

```

prevdepth  depth of the last line in the broken paragraph
prevgraf   number of lines in the broken paragraph
looseness  the actual looseness value in the broken paragraph
demerits   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`.

7.14.13.4 `tex.shipout`

```
tex.shipout(<number> n)
```

Ships out box number `n` to the output file, and clears the box register.

7.15 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	explanation
<code>kpse_init</code>	boolean	true	<code>false</code> totally disables <code>kpathsea</code> initialisation, and enables interpretation of the following numeric key-value pairs. (only ever unset this if you implement <i>all</i> file find callbacks!)
<code>shell_escape</code>	string	'f'	Use 'y' or 't' or 'l' to enable <code>\write 18</code> unconditionally, 'p' to enable the commands that are listed in <code>shell_escape_commands</code>
<code>shell_escape_commands</code>	string		Comma-separated list of command names that may be executed by <code>\write 18</code> even if <code>shell_escape</code> is set to 'p'. Do <i>not</i> use spaces around commas, separate any required command arguments by using a space, and use the ascii double quote (") for any needed argument or path quoting
<code>string_vacancies</code>	number	75000	cf. <code>web2c</code> docs
<code>pool_free</code>	number	5000	cf. <code>web2c</code> docs
<code>max_strings</code>	number	15000	cf. <code>web2c</code> docs
<code>strings_free</code>	number	100	cf. <code>web2c</code> docs
<code>nest_size</code>	number	50	cf. <code>web2c</code> docs
<code>max_in_open</code>	number	15	cf. <code>web2c</code> docs
<code>param_size</code>	number	60	cf. <code>web2c</code> docs
<code>save_size</code>	number	4000	cf. <code>web2c</code> docs
<code>stack_size</code>	number	300	cf. <code>web2c</code> docs
<code>dvi_buf_size</code>	number	16384	cf. <code>web2c</code> docs



<code>error_line</code>	number	79	cf. web2c docs
<code>half_error_line</code>	number	50	cf. web2c docs
<code>max_print_line</code>	number	79	cf. web2c docs
<code>hash_extra</code>	number	0	cf. web2c docs
<code>pk_dpi</code>	number	72	cf. web2c docs
<code>trace_file_names</code>	boolean	true	<code>false</code> disables T _E X's normal file open-close feedback (the assumption is that callbacks will take care of that)
<code>file_line_error</code>	boolean	false	do <code>file:line</code> style error messages
<code>halt_on_error</code>	boolean	false	abort run on the first encountered error
<code>formatname</code>	string		if no format name was given on the command line, this key will be tested first instead of simply quitting
<code>jobname</code>	string		if no input file name was given on the command line, this key will be tested first instead of simply giving up

Note: the numeric values that match web2c parameters are only used if `kpse_init` is explicitly set to `false`. In all other cases, the normal values from `texmf.cnf` are used.

7.16 The texio library

This library takes care of the low-level I/O interface.

7.16.1 Printing functions

7.16.1.1 `texio.write`

```
texio.write(<string> target, <string> s, ...)
texio.write(<string> s, ...)
```

Without the `target` argument, writes all given strings to the same location(s) T_EX 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.

7.16.1.2 `texio.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.



7.16.1.3 `texio.setescape`

You can disable ^^ escaping of control characters by passing a value of zero.

7.16.2 The `token` library

The current `token` library will be replaced by a new one that is more flexible and powerful. The transition takes place in steps. In version 0.80 we have `token` and in version 0.85 the old lib will be replaced completely. So if you use this new mechanism in production code you need to be aware of incompatible updates between 0.80 and 0.90. Because the related in- and output code will also be cleaned up and rewritten you should be aware of incompatible logging and error reporting too.

The old library presents tokens as triplets or numbers, the new library presents a userdata object. The old library used a callback to intercept tokens in the input but the new 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 LuaTeX is to provide a minimal set of tools and no solutions.

The current functions in the `token` namespace are given in the next table:

function	argument	result
<code>is_token</code>	<code>token</code>	checks if the given argument is a token userdata
<code>get_next</code>		returns the next token in the input
<code>scan_keyword</code>	<code>string</code>	returns true if the given keyword is gobbled
<code>scan_int</code>		returns a number
<code>scan_dimen</code>	<code>infinity</code> , <code>mu-units</code>	returns a number representing a dimension and or two numbers being the filler and order
<code>scan_glue</code>	<code>mu-units</code>	returns a glue spec node
<code>scan_toks</code>	<code>definer</code> , <code>expand</code>	returns a table of tokens token list (this can become a linked list in later releases)
<code>scan_code</code>	<code>bitset</code>	returns a character if its category is in the given bitset (representing catcodes)
<code>scan_string</code>		returns a string given between <code>{}</code> , as <code>\macro</code> or as sequence of characters with catcode 11 or 12
<code>scan_word</code>		returns a sequence of characters with catcode 11 or 12 as string
<code>scan_csname</code>		returns <code>foo</code> after scanning <code>\foo</code>
<code>set_macro</code>	see below	assign a macro
<code>create</code>		returns a userdata token object of the given control sequence name (or character); this interface can change

The scanners can be considered stable apart from the one scanning for a token. This is because futures releases can return a linked list instead of a table (as with nodes). 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:

```
\directlua{token.scan_string(){foo}  foo      full expansion
\directlua{token.scan_string()}foo    foo      letters and others
\directlua{token.scan_string()\foo}  foo-bar  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 as 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.

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:

<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>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

The numbers that represent a catcode are the same as in $\text{T}_{\text{E}}\text{X}$ itself, so using this information assumes that you know a bit about $\text{T}_{\text{E}}\text{X}$'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.

More interesting are the scanners. 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 when it's 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).

When the interface is stable and has replaced the old one completely we will add some more information here. By that time the internals have been cleaned up a bit more so we know then what will stay and go. A positive side effect of this transition is that we can simplify the input part because we no longer need to intercept using callbacks.



The `set_macro` function can get upto 4 arguments:

```
setmacro("csname", "content")
setmacro("csname", "content", "global")
setmacro("csname")
```

You can pass a `catcodetable` identifier as first argument:

```
setmacro(catcodetable, "csname", "content")
setmacro(catcodetable, "csname", "content", "global")
setmacro(catcodetable, "csname")
```

The results are like:

```
\def\csname{content}
\gdef\csname{content}
\def\csname{}
```

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.





8 Modifications

8.1 The merged engines

8.1.1 The need for change

The first version of Lua \TeX only had a few extra primitives and it was largely the same as pdf \TeX . Then we merged substantial parts of Aleph into the code and got more primitives. When we got more stable the decision was made to clean up the rather hybrid nature of the program. This means that some primitives have been promoted to core primitives, often with a different name, and that others were removed. This made it possible to start cleaning up the code base. We will describe most in following paragraphs.

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.

8.1.2 Changes from \TeX 3.1415926

Of course it all starts with traditional \TeX . Even if we started with pdf \TeX , most still comes from the original. But we divert a bit.

- The current code base is written in C, not Pascal. We use cweb when possible.
- See chapter 3 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.
- There is no pool file, all strings are embedded during compilation.
- The specifier `plus 1 filllll` does not generate an error. The extra 'l' is simply typeset.
- The upper limit to `\endlinechar` and `\newlinechar` is 127.
- 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 more efficient (pdf) code is generated.

8.1.3 Changes from $\epsilon\text{-}\TeX$ 2.2

Being the de facto standard extension of course we provide the $\epsilon\text{-}\TeX$ functionality, but with a few small adaptations.

- The $\epsilon\text{-}\TeX$ functionality is always present and enabled so the prepended asterisk or `-etex` switch for `ini \TeX` is not needed.
- The $\TeX\text{XeT}$ extension is not present, so the primitives `\TeXXeTstate`, `\beginR`, `\beginL`, `\endR` and `\endL` are missing.
- Some of the tracing information that is output by $\epsilon\text{-}\TeX$'s `\tracingassigns` and `\tracingrestores` is not there.



- Register management in Lua \TeX uses the Aleph model, so the maximum value is 65535 and the implementation uses a flat array instead of the mixed flat&sparse model from ϵ - \TeX .
- When `kpathsea` is used to find files, Lua \TeX uses the `ofm` file format to search for font metrics. In turn, this means that Lua \TeX looks at the `OFM FONTS` configuration variable (like Omega and Aleph) instead of `TFM FONTS` (like \TeX and pdf \TeX). Likewise for virtual fonts (Lua \TeX uses the variable `OVFFONTS` instead of `VFFONTS`).

8.1.4 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 Lua \TeX functionality which in turn triggered renaming primitives.

- The (experimental) support for snap nodes has been removed, because it is much more natural to build this functionality on top of node processing and attributes. The associated primitives that are now gone are: `\pdfsnaprefpoint`, `\pdfsnapy`, and `\pdfsnapycomp`.
- The (experimental) support for specialized spacing around nodes has also been removed. The associated primitives that are now gone are: `\pdfadjustinterwordglue`, `\pdfprependkern`, and `\pdfappendkern`, as well as the five supporting primitives `\knbscode`, `\stbscode`, `\shbscode`, `\knbcode`, and `\knaccode`.
- A number of ‘pdf \TeX primitives’ have been removed as they can be implemented using Lua: `\pdfelapsetime`, `\pdfescapehex`, `\pdfescapename`, `\pdfescapestring`, `\pdffiledump`, `\pdffilemoddate`, `\pdffilesize`, `\pdfforcepagebox`, `\pdflastmatch`, `\pdfmatch`, `\pdfmd-fivesum`, `\pdfmovechars`, `\pdfoptionalwaysusepdfpagebox`, `\pdfoptionpdfinclusion-errorlevel`, `\pdfresettimer`, `\pdfshellescape`, `\pdfstrcmp` and `\pdfunescapehex`
- The version related primitives `\pdftexbanner`, `\pdftexversion` and `\pdftexrevision` are no longer present as there is no longer a strict relationship with pdf \TeX development.
- The experimental snapper mechanism has been removed and therefore also the primitives: `\pdfignoreddimen`, `\pdffirstlineheight`, `\pdfeachlineheight`, `\pdfeachlinedepth` and `\pdflastlinedepth`
- The experimental primitives `\primitive`, `\ifprimitive`, `\ifabsnum` and `\ifabsdim` are promoted to core primitives. The `\pdf*` prefixed originals are not available.
- The png transparency fix from 1.40.6 is not applied as high-level support is pending.
- Two extra token lists are provided, `\pdfxformresources` and `\pdfxformattr`, as an alternative to `\pdfxform` keywords.
- The current version of Lua \TeX no longer replaces and/or merges fonts in embedded pdf files with fonts of the enveloping pdf document. This regression may be temporary, depending on how the rewritten font backend will look like.
- The primitives `\pdfpagewidth` and `\pdfpageheight` have been removed because `\pagewidth` and `\pageheight` have that purpose.
- The primitives `\pdfnormaldeviate`, `\pdfuniformdeviate`, `\pdfsetrandomseed` and `\pdfrandomseed` have been promoted to core primitives without `pdf` prefix so the original commands are no longer recognized.
- The primitives `\ifincsname`, `\expanded` and `\quitvmode` are now core primitives.
- As the hz and protrusion mechanism are part of the core the related primitives `\lpcode`, `\rpcode`, `\efcode`, `\leftmarginkern`, `\rightmarginkern` are promoted to core primitives.



The two commands `\protrudechars` and `\adjustspacing` replace their prefixed with `\pdf` originals.

- 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.
- The `\tagcode` primitive is promoted to core primitive.
- The `\letterspacefont` feature is now part of the core but will not be changed (improved). We just provide it for legacy use.
- The `\pdfnoligatures` primitive is now `\ignoreligaturesinfont`.
- The `\pdfcopyfont` primitive is now `\copyfont`.
- The `\pdffontexpand` primitive is now `\expandglyphsinfont`.
- Because position tracking is also available in dvi mode the `\savepos`, `\lastxpos` and `\lastypos` commands now replace their `pdf` prefixed originals.
- The introspective primitives `\pdflastximagecolordepth` and `\pdfximagebbox` have been removed. One can use external applications to determine these properties or use the built-in `img` library.
- The initializers `\pdfoutput` has been replaced by `\outputmode` and `\pdfdraftmode` is now `\draftmode`.
- The pixel multiplier dimension `\pdfpxdimen` loses its prefix and is now called `\pxdimen`.
- An extra `\pdfimageaddfilename` option has been added that can be used to block writing the filename to the pdf file.

One change involves the so called xforms and ximages. In pdf_TE_X these are implemented as so called whatsits. But contrary to other whatsits they have dimensions that need to be taken into account when for instance calculating optimal line breaks. In Lua_TE_X these are now promoted to normal nodes, which simplifies code that needs those dimensions.

Another reason for promotion is that these are useful concepts. Backends can provide the ability to use content that has been rendered in several places, and images are also common. For that reason we also changed the names:

new name	old name
<code>\saveboxresource</code>	<code>\pdfxform</code>
<code>\saveimageresource</code>	<code>\pdfximage</code>
<code>\useboxresource</code>	<code>\pdfrefxform</code>
<code>\useimageresource</code>	<code>\pdfrefximage</code>
<code>\lastsavedboxresourceindex</code>	<code>\pdflastxform</code>
<code>\lastsavedimageresourceindex</code>	<code>\pdflastximage</code>
<code>\lastsavedimageresourcepages</code>	<code>\pdflastximagepages</code>

There are a few `\pdf...` primitives that relate to this but these are typical backend specific ones. The index that gets returned is to be considered as ‘just a number’ and although it still has the same meaning (object related) as before, you should not depend on that.

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 boundary node to make the next node being ignored. When the boundary 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:

```
\boundary1\llap{!\quad}\llap{«Who needs protrusion?»}
```

8.1.5 Changes from ALEPH RC4

Because we wanted proper directional typesetting the Aleph mechanisms looked most attractive. These are rather close to the ones provided by Omega, so what we say next applies to both these programs.

- The extended 16-bit math primitives (`\omathcode` etc.) have been removed.
- The OCP processing is no longer supported at all. As a consequence, the following primitives have been removed:
`\ocp`, `\externalocp`, `\ocplist`, `\pushocplist`, `\popocplist`, `\clearocplists`, `\addbeforeocplist`, `\addafterocplist`, `\removebeforeocplist`, `\removeafterocplist` and `\ocptracelevel`
- LuaTeX only understands 4 of the 16 direction specifiers of Aleph: `TLT` (latin), `TRT` (arabic), `RTT` (cjk), `LTL` (mongolian). All other direction specifiers generate an error.
- The input translations from Aleph are not implemented, the related primitives are not available:
`\DefaultInputMode`, `\noDefaultInputMode`, `\noInputMode`, `\InputMode`, `\DefaultOutputMode`, `\noDefaultOutputMode`, `\noOutputMode`, `\OutputMode`, `\DefaultInputTranslation`, `\noDefaultInputTranslation`, `\noInputTranslation`, `\InputTranslation`, `\DefaultOutputTranslation`, `\noDefaultOutputTranslation`, `\noOutputTranslation` and `\OutputTranslation`
- Several bugs have been fixed. The `\hoffset` bug when `\pagedir TRT` is gone, removing the need for an explicit fix to `\hoffset`. Also bug causing `\fam` to fail for family numbers above 15 is fixed. A fair amount of other minor bugs are fixed as well, most of these related to `\tracingcommands` output.
- The scanner for direction specifications now allows an optional space after the direction is completely parsed.
- The `^^` notation has been extended: after `^^^^` four hexadecimal characters are expected and after `^^^^^^` six hexadecimal characters have to be given. The original TeX 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.
- Several mechanisms that need to be right-to-left aware have been improved. For instance placement of formula numbers.
- The page dimension related primitives `\pagewidth` and `\pageheight` have been promoted to core primitives.
- The primitives `\charwd`, `\charht`, `\chardp` and `\charit` have been removed as we have the ε -T_EX variants `\fontchar*`.
- The two dimension registers `\pagerightoffset` and `\pagebottomoffset` are now core primitives.
- The direction related primitives `\pagedir`, `\bodydir`, `\pardir`, `\textdir`, `\mathdir` and `\boxdir` are now core primitives.
- The primitive `\pdftracingfonts` is now `\tracingfonts` as it doesn't relate to the backend.
- The experimental primitive `\pdfinsertht` is kept as `\insertht`.
- The promotion of primitives to core primitives as well as the removal of all others mean that the initialization namespace `aleph` is gone.

8.1.6 Changes from standard WEB2C

The compilation framework is web2c and we keep using that but without the Pascal to C step. This framework also provides some common features that deal with reading bytes from files and locating files in tds. This is what we do different:

- There is no mltex support.
- There is no enctex support.
- The following commandline switches are silently ignored, even in non-Lua mode: `-8bit`, `-translate-file`, `-mltex`, `-enc` and `-etex`.
- The `\openout` whatsits are not written to the log file.
- Some of the so-called web2c extensions are hard to set up in non-kpse mode because `texmf.cnf` is not read: `shell-escape` is off (but that is not a problem because of Lua's `os.execute`), and the paranoia checks on `openin` and `openout` do not happen (however, it is easy for a Lua script to do this itself by overloading `io.open`).
- The 'E' option does not do anything useful.

8.1.7 The backend primitives `\pdf` *

In a previous section we mentioned that some pdfT_EX primitives were removed and others promoted to core LuaT_EX primitives. That is only part of the story. In order to separate the backend specific primitives in the code these commands are now replaced by only a few. In traditional T_EX we only had the dvi backend but now we have two: dvi and pdf. Additional functionality is implemented as 'extensions' in T_EXspeak. By separating more strictly we are able to keep the core (fontend) clean and stable. If for some reason an extra backend option is needed, it can be implemented without touching the core. The three pdf backend related primitives are

`\pdfextension command [specification]`



```
\pdfvariable name
\pdffeedback name
```

An extension triggers further parsing, depending on the command given. A variable is a (kind of) register and can be read and written, while a feedback is reporting something (as it comes from the backend it's normally a sequence of tokens).

In order for LuaT_EX to be more than just T_EX you need to enable primitives. That has already been the case right from the start. If you want the traditional pdfT_EX primitives (for as far their functionality is still around) you now can do this:

```
\protected\def\pdfliteral          {\pdfextension literal}
\protected\def\pdfcolorstack       {\pdfextension colorstack}
\protected\def\pdfsetmatrix        {\pdfextension setmatrix}
\protected\def\pdfsave             {\pdfextension save\relax}
\protected\def\pdfrestore         {\pdfextension restore\relax}
\protected\def\pdfobj             {\pdfextension obj }
\protected\def\pdfrefobj          {\pdfextension refobj }
\protected\def\pdfannot           {\pdfextension annot }
\protected\def\pdfstartlink       {\pdfextension startlink }
\protected\def\pdfendlink         {\pdfextension endlink\relax}
\protected\def\pdfoutline         {\pdfextension outline }
\protected\def\pdfdest            {\pdfextension dest }
\protected\def\pdfthread          {\pdfextension thread }
\protected\def\pdfstartthread     {\pdfextension startthread }
\protected\def\pdfendthread       {\pdfextension endthread\relax}
\protected\def\pdfinfo            {\pdfextension info }
\protected\def\pdfcatalog         {\pdfextension catalog }
\protected\def\pdfnames           {\pdfextension names }
\protected\def\pdfincludechars    {\pdfextension includechars }
\protected\def\pdffontattr        {\pdfextension fontattr }
\protected\def\pdfmapfile         {\pdfextension mapfile }
\protected\def\pdfmapline         {\pdfextension mapline }
\protected\def\pdftrailer         {\pdfextension trailer }
\protected\def\pdfglyphtounicode  {\pdfextension glyphtounicode }
\protected\def\pdfsuppressoptionalinfo{\pdfextension suppressoptionalinfo }
\protected\def\pdfinvoid          {\pdfextension invoid }
```

The introspective primitives can be defined as:

```
\def\pdfTeXversion  {\numexpr\pdffeedback version\relax}
\def\pdfTeXrevision {\pdffeedback revision}
\def\pdfLastlink    {\numexpr\pdffeedback lastlink\relax}
\def\pdfRetval      {\numexpr\pdffeedback retval\relax}
\def\pdfLastobj     {\numexpr\pdffeedback lastobj\relax}
\def\pdfLastannot   {\numexpr\pdffeedback lastannot\relax}
\def\pdfXformname   {\numexpr\pdffeedback xformname\relax}
\def\pdfCreationdate {\pdffeedback creationdate}
```



```

\def\pdffontname      {\numexpr\pdffeedback fontname\relax}
\def\pdffontobjnum   {\numexpr\pdffeedback fontobjnum\relax}
\def\pdffontsize     {\dimexpr\pdffeedback fontsize\relax}
\def\pdfpageref      {\numexpr\pdffeedback pageref\relax}
\def\pdfcolorstackinit  {\pdffeedback colorstackinit}

```

The configuration related registers have become:

```

\edef\pdfcompresslevel      {\pdfvariable compresslevel}
\edef\pdfobjcompresslevel   {\pdfvariable objcompresslevel}
\edef\pdfdecimaldigits     {\pdfvariable decimaldigits}
\edef\pdfgamma              {\pdfvariable gamma}
\edef\pdfimageresolution   {\pdfvariable imageresolution}
\edef\pdfimageapplygamma   {\pdfvariable imageapplygamma}
\edef\pdfimagegamma        {\pdfvariable imagegamma}
\edef\pdfimagehicolor      {\pdfvariable imagehicolor}
\edef\pdfimageaddfilename  {\pdfvariable imageaddfilename}
\edef\pdfpkresolution      {\pdfvariable pkresolution}
\edef\pdfpkfixeddpi        {\pdfvariable pkfixeddpi}
\edef\pdfinclusioncopyfonts  {\pdfvariable inclusioncopyfonts}
\edef\pdfinclusionerrorlevel {\pdfvariable inclusionerrorlevel}
\edef\pdfignoreunknownimages {\pdfvariable ignoreunknownimages}
\edef\pdfgentounicode      {\pdfvariable gentounicode}
\edef\pdfpagebox           {\pdfvariable pagebox}
\edef\pdfminorversion      {\pdfvariable minorversion}
\edef\pdfuniqueresname     {\pdfvariable uniqueresname}

\edef\pdfhorigin           {\pdfvariable horigin}
\edef\pdfvorigin           {\pdfvariable vorigin}
\edef\pdflinkmargin        {\pdfvariable linkmargin}
\edef\pdfdestmargin        {\pdfvariable destmargin}
\edef\pdfthreadmargin      {\pdfvariable threadmargin}
\edef\pdfxformmargin       {\pdfvariable xformmargin}

\edef\pdfpagesattr         {\pdfvariable pagesattr}
\edef\pdfpageattr          {\pdfvariable pageattr}
\edef\pdfpageresources     {\pdfvariable pageresources}
\edef\pdfxformattr         {\pdfvariable xformattr}
\edef\pdfxformresources    {\pdfvariable xformresources}
\edef\pdfpkmode            {\pdfvariable pkmode}

```

The variables are internal ones, so they are anonymous. When you ask for the meaning of a few previously defined ones:

```

\meaning\pdfhorigin
\meaning\pdfcompresslevel
\meaning\pdfpageattr

```



you will get:

```
macro:->[internal backend dimension]
macro:->[internal backend integer]
macro:->[internal backend tokenlist]
```

The `\edef` can also be an `\def` but it's a bit more efficient to expand the lookup related register beforehand. After that you can adapt the defaults; these are:

```
\pdfcompresslevel          9
\pdfobjcompresslevel      1 % used: (0,9)
\pdfdecimaldigits        4 % used: (3,6)
\pdfgamma                 1000
\pdfimageresolution       71
\pdfimageapplygamma       0
\pdfimagegamma            2200
\pdfimagehicolor          1
\pdfimageaddfilename      1
\pdfpkresolution          72
\pdfpkfixeddpi            0
\pdfinclusioncopyfonts     0
\pdfinclusionerrorlevel    0
\pdfignoreunknownimages  0
\pdfgentounicode         0
\pdfpagebox               0
\pdfminorversion          4
\pdfuniquestname         0

\pdfhorigin               1in
\pdfvorigin               1in
\pdflinkmargin            0pt
\pdfdestmargin           0pt
\pdfthreadmargin         0pt
\pdfxformmargin          0pt
```

If you also want some backward compatibility, you can add:

```
\let\pdfpagewidth         \pagewidth
\let\pdfpageheight       \pageheight

\let\pdfadjustspacing    \adjustspacing
\let\pdfprotrudechars    \protrudechars
\let\pdfnoligatures      \ignoreligaturesinfont
\let\pdffontexpand       \expandglyphsinfont
\let\pdfcopyfont         \copyfont

\let\pdfxform            \saveboxresource
\let\pdflastxform       \lastsavedboxresourceindex
```



```

\let\pdfrefxform      \useboxresource

\let\pdfximage        \saveimageresource
\let\pdflastximage    \lastsavedimageresourceindex
\let\pdflastximagepages\lastsavedimageresourcepages
\let\pdfrefximage     \useimageresource

\let\pdfsavepos       \savepos
\let\pdflastxpos      \lastxpos
\let\pdflastypos      \lastypos

\let\pdfoutput        \outputmode
\let\pdfdraftmode     \draftmode

\let\pdfpxdimen       \pxdimen

\let\pdfinsertht      \insertht

\let\pdfnormaldeviate \normaldeviate
\let\pdfuniformdeviate \uniformdeviate
\let\pdfsetrandomseed \setrandomseed
\let\pdfrandomseed    \randomseed

\let\pdfprimitive     \primitive
\let\ifpdfprimitive   \ifprimitive

\let\ifpdfabsnum      \ifabsnum
\let\ifpdfabsdim      \ifabsdim

```

And even:

```

\newdimen\pdfeachlineheight
\newdimen\pdfeachlinedepth
\newdimen\pdflastlinedepth
\newdimen\pdffirstlineheight
\newdimen\pdfignoreddimen

```

The backend is derived from pdf_TE_X so the same syntax applies. However, the `outline` command accepts a `objnum` followed by a number. No checking takes place so when this is used it had better be a valid (flushed) object.

In order to be (more or less) compatible with pdf_TE_X we also support the option to suppress some info:

```

\pdfvariable suppressoptionalinfo \numexpr
  0
+ 1 % PTEX.FullBanner
+ 2 % PTEX.FileName

```



```

+ 4 % PTEX.PageNumber
+ 8 % PTEX.InfoDict
+ 16 % Creator
+ 32 % CreationDate
+ 64 % ModDate
+ 128 % Producer
+ 256 % Trapped
+ 512 % ID
\relax

```

In addition you can overload the trailer id, but we don't do any checking on validity, so you have to pass a valid array. The following is like the ones normally generated by the engine:

```

\pdfvariable trailerid {[
  <FA052949448907805BA83C1E78896398>
  <FA052949448907805BA83C1E78896398>
]}

```

So, you even need to include the brackets!

8.2 Implementation notes

8.2.1 Memory allocation

The single internal memory heap that traditional $\text{T}_{\text{E}}\text{X}$ used for tokens and nodes is split into two separate arrays. Each of these will grow dynamically when needed.

The `texmf.cnf` settings related to main memory are no longer used (these are: `main_memory`, `mem_bot`, `extra_mem_top` and `extra_mem_bot`). 'Out of main memory' errors can still occur, but the limiting factor is now the amount of RAM in your system, not a predefined limit.

Also, the memory (de)allocation routines for nodes are completely rewritten. The relevant code now lives in the C file `texnode.c`, and basically uses a dozen or so 'avail' lists instead of a doubly-linked model. An extra function layer is added so that the code can ask for nodes by type instead of directly requisitioning a certain amount of memory words.

Because of the split into two arrays and the resulting differences in the data structures, some of the macros have been duplicated. For instance, there are now `vlink` and `vinfo` as well as `token_link` and `token_info`. All access to the variable memory array is now hidden behind a macro called `vmem`.

The implementation of the growth of two arrays (via reallocation) introduces a potential pitfall: the memory arrays should never be used as the left hand side of a statement that can modify the array in question.

The input line buffer and pool size are now also reallocated when needed, and the `texmf.cnf` settings `buf_size` and `pool_size` are silently ignored.



8.2.2 Sparse arrays

The `\mathcode`, `\delcode`, `\catcode`, `\sfcode`, `\lccode` and `\uccode` 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.

The `\catcode`, `\sfcode`, `\lccode` and `\uccode` assignments do not yet show up when using the etex tracing routines `\tracingassigns` and `\tracingrestores` (code simply not written yet).

A side-effect of the current implementation is that `\global` is now more expensive in terms of processing than non-global assignments.

See `mathcodes.c` and `textcodes.c` if you are interested in the details.

Also, the glyph ids within a font are now managed by means of a sparse array and glyph ids can go up to index $2^{21} - 1$.

8.2.3 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.

8.2.4 Compressed format

The format is passed through zlib, allowing it to shrink to roughly half of the size it would have had in uncompressed form. This takes a bit more cpu cycles but much less disk io, so it should still be faster.

8.2.5 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).



