

# MKII – MKIV

## Abstract

This article is the first in a series about ConT<sub>E</sub>Xt and LuaT<sub>E</sub>X. For those who use ConT<sub>E</sub>Xt it is a progress report of the development process and the choices that are being made. For those not using ConT<sub>E</sub>Xt it gives some insight in what LuaT<sub>E</sub>X is about.

## From Mark II to Mark IV

Sometime in 2005 the development of LuaT<sub>E</sub>X started, a further development of PDFT<sub>E</sub>X and a precursor to PDFT<sub>E</sub>X version 2. This T<sub>E</sub>X variant will provide:

- 21–32 bit internals plus a code cleanup
- flexible support for OpenType fonts
- an internal UTF data flow
- the bidirectional typesetting of ALEPH
- LUA callbacks to the most relevant T<sub>E</sub>X internals
- some extensions to T<sub>E</sub>X (for instance math)
- an efficient way to communicate with MetaPost

In the tradition of T<sub>E</sub>X this successor will be downward compatible in most essential parts and in the end, there is still PDFT<sub>E</sub>X version 1 as fall back.

In the mean time we have seen another unicode variant show up, X<sub>Y</sub>T<sub>E</sub>X which is under active development, uses external libraries, provides access to the fonts on the operating system, etc.

From the beginning, ConT<sub>E</sub>Xt always worked with all engines. This was achieved by conditional code blocks: depending on what engine was used, different code was put in the format and/or used at runtime. Users normally were unaware of this. Examples of engines are  $\epsilon$ -T<sub>E</sub>X, ALEPH, and X<sub>Y</sub>T<sub>E</sub>X. Because nowadays all engines provide the  $\epsilon$ -T<sub>E</sub>X features, in August 2006 we decided to consider those features to be present and drop providing the standard T<sub>E</sub>X compatible variants. This is a small effort because all code that is sensitive for optimization already has  $\epsilon$ -T<sub>E</sub>X code branches for many years.

However, with the arrival of LuaT<sub>E</sub>X, we need a more drastic approach. Quite some existing code can go away and will be replaced by different solutions. Where T<sub>E</sub>X code ends up in the format file, along with its state, LUA code will be initiated at run time, after a LUA instance is started. ConT<sub>E</sub>Xt reserves its own instance of LUA.

Most of this will go unnoticed for the users because the user interface will not change. For developers however, we need to provide a mechanism to deal with these issues. This is why, for the first time in ConT<sub>E</sub>Xt's history we will officially use a kind of version tag. When we changed the low level interface from Dutch to English we jokingly talked of version 2. So, it makes sense to follow this lead.

- ConT<sub>E</sub>Xt Mark I At that moment we still had a low level Dutch interface, invisible for users but not for developers.
- ConT<sub>E</sub>Xt Mark II We now have a low level English interface, which (as we indeed saw happen) triggers more development by users.
- ConT<sub>E</sub>Xt Mark IV This is the next generation of ConT<sub>E</sub>Xt, with parts re-implemented. It's an at some points drastic system overhaul.

Keep in mind that the functionality does not change, although in some places, for instance fonts, Mark IV may provide additional functionality. The reason why most users will not notice the difference (maybe apart from performance and convenience) is that at the user interface level nothing changes (most of it deals with typesetting, not with low level details).

The hole in the numbering permits us to provide a Mark III version as well. Once X<sub>Y</sub>TeX is stable, we may use that slot for X<sub>Y</sub>TeX specific implementations.

As per August 2006 the banner is adapted to this distinction:

```
... ver: 2006.09.06 22:46 MK II  fmt: 2006.9.6  ...
... ver: 2006.09.06 22:47 MK IV  fmt: 2006.9.6  ...
```

This numbering system is reflected at the file level in such a way that we can keep developing the way we do, i.e. no files all over the place, in subdirectories, etc.

Most of the system's core files are not affected, but some may be, like those dealing with fonts, input- and output encodings, file handling, etc. Those files may come with different suffixes:

- `somefile.tex`: the main file, implementing the interface and common code
- `somefile.mkii`: mostly existing code, suitable for good old TeX ( $\epsilon$ -TeX, PDFTeX, ALEPH).
- `somefile.mkiv`: code optimized for use with L<sup>A</sup>TeX, which could follow completely different approaches
- `somefile.lua`: LUA code, loaded at format generation time and/or runtime

As said, some day `somefile.mkiii` code may show up. Which variant is loaded is determined automatically at format generation time as well as at run time.

## How LUA fits in

### introduction

Here I will discuss a few of the experiments that drove the development of L<sup>A</sup>TeX. It describes the state of affairs around the time that we were preparing for TUG 2006. This development was pretty demanding for Taco and me but also much fun. We were in a kind of permanent Skype chat session, with binaries flowing in one direction and TeX and LUA code the other way. By gradually replacing (even critical) components of ConTeXt we had a real test bed and torture tests helped us to explore and debug at the same time. Because Taco uses LINUX as platform and I mostly use MS WINDOWS, we could investigate platform dependent issues conveniently. While reading this text, keep in mind that this is just the beginning of the game.

I will not provide sample code here. When possible, the Mark IV code transparently replaces Mark II code and users will seldom notices that something happens in different way. Of course the potential is there and future extensions may be unique to Mark IV.

### compatibility

The first experiments, already conducted with the experimental versions involved runtime conversion of one type of input into another. An example of this is the (TI) calculator math input handler that converts a rather natural math sequence into TeX and feeds that back into TeX. This mechanism eventually will evolve into a configurable math input handler. Such applications are unique to Mark IV code and will not be backported to Mark II. The question is where downward compatibility will become a problem. We don't expect many problems, apart from occasional bugs that result from splitting the code base, mostly because new features will not affect older functionality. Because we have to reorganize the code base a bit, we also use this opportunity to start making a variant of ConTeXt which consists of building blocks: METATEX. This is less interesting for the average user, but may be

of interest for those using ConT<sub>E</sub>Xt in workflows where only part of the functionality is needed.

### metapost

Of course, when I experiment with such new things, I cannot let MetaPost leave untouched. And so, in the early stage of L<sup>A</sup>T<sub>E</sub>X development I decided to play with two MetaPost related features: conversion and runtime processing.

Conversion from MetaPost output to PDF is currently done in pure T<sub>E</sub>X code. Apart from convenience, this has the advantage that we can let T<sub>E</sub>X take care of font inclusions. The tricky part of this conversion is that MetaPost output has some weird aspects, like DVIPS specific linewidth snapping. Another nasty element in the conversion is that we need to transform paths when pens are used. Anyhow, the converter has reached a rather stable state by now.

One of the ideas with MetaPost version 1<sup>+</sup> is that we will have an alternative output mode. In the perspective of L<sup>A</sup>T<sub>E</sub>X it makes sense to have a LUA output mode. Whatever converter we use, it needs to deal with METAFUN specials. These are responsible for special features like transparency, graphic inclusion, shading, and more. Currently we misuse colors to signal such features, but the new pre/post path hooks permit more advanced implementations. Experimenting with such new features is easier in LUA than in T<sub>E</sub>X.

The Mark IV converter is a multi-pass converter. First we clean up the MetaPost output, next we convert the PostScript code into LUA calls. We assume that this LUA code eventually can be output directly from MetaPost. We then evaluate this converted LUA blob, which results in T<sub>E</sub>X commands. Think of:

```
1.2 setlinejoin
```

turned into:

```
mp.setlinejoin(1.2)
```

becoming:

```
\PDFcode{1.2 j}
```

which is, when the PDF<sub>T</sub>E<sub>X</sub> driver is active, equivalent to:

```
\pdfliteral{1.2 j}
```

Of course, when paths are involved, more things happen behind the scenes, but in the end an mp.path enters the LUA machinery.




Figure 1. converter test figure

prologues/mpprocset	1/0	1/1	2/0/1
Mark II	8.5 ( 5.7)	8.0 (5.5)	8.8 8.5
Mark IV	16.1 (10.6)	7.2 (4.5)	16.3 7.4

**Table 1.** converter speed benchmark data

When the Mark IV converter reached a stable state, tests demonstrated then the code was upto 20% slower than the pure  $\text{\TeX}$  alternative on average graphics, and but faster when many complex path transformations (due to penshapes) need to be done. This slowdown was due to the cleanup (using expressions) and intermediate conversion. Because Taco develops  $\text{\LaTeX}$  as well as maintains and extends MetaPost, we conducted experiments that combine features of these programs. As a result of this, shortcuts found their way into the MetaPost output.

Cleaning up the MetaPost output using LUA expressions takes relatively much time. However, starting with version 0.970 MetaPost uses a preamble, which permits not only short commands, but also gets rid of the weird linewidth and filldraw related PostScript constructs. The moderately complex graphic that we use for testing (figure 1) takes over 16 seconds when converted 250 times. When we enable shortcuts we can avoid part of the cleanup and runtime goes down to under 7.5 seconds. This is significantly faster than the Mark II code. We did experiments with simulated LUA output from MetaPost and then the Mark IV converter really flies. The benchmark data are shown in table 1. The values on Taco's system are given between parenthesis.

The main reason for the huge difference in the Mark IV times is that we do a rigorous cleanup of the older MetaPost output in order avoid messy the messy (but fast) code that we use in the Mark II converter. Think of:

```
0 0.5 dtransform truncate idtransform setlinewidth pop
closepath gsave fill grestore stroke
```

In the Mark II converter, we push every number or keyword on a stack and use keywords as trigger points. In the Mark IV code we convert the stack based PostScript calls to LUA function calls. Lines as shown are converted to single calls first. When prologues is set to 2, such line no longer show up and are replaced by simple calls accompanied by definitions in the preamble. Not only that, instead of verbose keywords, one or two character shortcuts are used. This means that the Mark II code can be faster when procsets are used because shorter strings end up in the stack and comparison happens faster. On the other hand, when no procsets are used, the runtime is longer because of the larger preamble.

Because the converter is used outside  $\text{Con}\text{\TeX}$ t as well, we support all combinations in order not to get error messages, but the converter is supposed to work with the following settings:

```
prologues := 1 ;
mpprocset := 1 ;
```

We don't need to set prologues to 2 (font encodings in file) or 3 (also font resources in file). So, in the end, the comparison in speed comes down to 8.0 seconds for Mark II code and 7.2 seconds for the Mark IV code when using the latest greatest MetaPost. When we simulate LUA output from MetaPost, we end up with 4.2 seconds runtime and when MetaPost could produce the converter's  $\text{\TeX}$  commands, we need only 0.3 seconds for embedding the 250 instances. This includes  $\text{\TeX}$  taking care of handling the specials, some of which demand building moderately complex PDF data structures.

But, conversion is not the only factor in convenient MetaPost usage. First of all, runtime MetaPost processing takes time. The actual time spent on handling

embedded MetaPost graphics is also dependent on the speed of starting up MetaPost, which in turn depends on the size of the  $\TeX$  trees used: the bigger these are, the more time  $\text{kpse}$  spends on loading the  $\text{ls-R}$  databases. Eventually this bottleneck may go away when we have MetaPost as a library. (In  $\text{Con}\TeX$ t one can also run MetaPost between runs. Which method is faster, depends on the amount and complexity of the graphics.)

Another factor in dealing with MetaPost, is the usage of text in a graphic ( $\text{btex}$ ,  $\text{ttext}$ , etc.). Taco Hoekwater, Fabrice Popineau and I did some experiments with a persistent MetaPost session in the background in order to simulate a library. The results look very promising: the overhead of embedded MetaPost graphics goes to nearly zero, especially when we also let the parent  $\TeX$  job handle the typesetting of texts. A side effect of these experiments was a new mechanism in  $\text{Con}\TeX$ t (and  $\text{META}\text{FUN}$ ) where  $\TeX$  did all typesetting of labels, and MetaPost only worked with an abstract representation of the result. This way we can completely avoid nested  $\TeX$  runs (the ones triggered by MetaPost). This also works ok in Mark II mode.

Using a persistent MetaPost run and piping data into it is not the final solution if only because the terminal log becomes messed up too much, and also because intercepting errors is real messy. In the end we need a proper library approach, but the experiments demonstrated that we needed to go this way: handling hundreds of complex graphics that hold typeset paragraphs (being slanted and rotated and more by MetaPost), took mere seconds compared to minutes when using independent MetaPost runs for each job.

#### characters

Because  $\text{LUA}\TeX$  is UTF based, we need a different way to deal with input encoding. For this purpose there are callbacks that intercept the input and convert it as needed. For context this means that the regime related modules get a LUA based counterparts. As a prelude to advanced character manipulations, we already load extensive unicode and conversion tables, with the benefit of being able to handle case handling with LUA.

The character tables are derived from unicode tables and Mark II  $\text{Con}\TeX$ t data files and generated using  $\text{MTX}\text{TOOLS}$ . The main character table is pretty large, and this made us experiment a bit with efficiency. It was in this stage that we realized that it made sense to use precompiled LUA code (using  $\text{luac}$ ). During format generation we let  $\text{Con}\TeX$ t keep track of used LUA files and compiled them on the fly. For a production run, the compiled files were loaded instead.

Because at that stage  $\text{LUA}\TeX$  was already a merge between  $\text{PDF}\TeX$  and  $\text{ALEPH}$ , we had to deal with pretty large format files. About that moment the  $\text{Con}\TeX$ t format with the english user interface amounted to:

date	luatex	pdftex	xetex	aleph
2006-09-18	9 552 042	7 068 643	8 374 996	7 942 044

One reason for the large size of the format file is that the memory footprint of a 32 bit  $\TeX$  is larger than that of good old  $\TeX$ , even with some of the clever memory allocation techniques as used in  $\text{LUA}\TeX$ . After some experiments where size and speed were measured Taco decided to compress the format using a level 3 ZIP compression. This brilliant move lead to the following size:

date	luatex	pdftex	xetex	aleph
2006-10-23	3 135 568	7 095 775	8 405 764	7 973 940

The first zipped versions were smaller (around 2.3 meg), but in the meantime we moved the LUA code into the format and the character related tables take some space.

## debugging

In the process of experimenting with callbacks I played a bit with handling  $\TeX$  error information. An option is to generate an HTML page instead of spitting out the usual blob of into on the terminal. In figure 2 and figure 3 you can see an example of this.

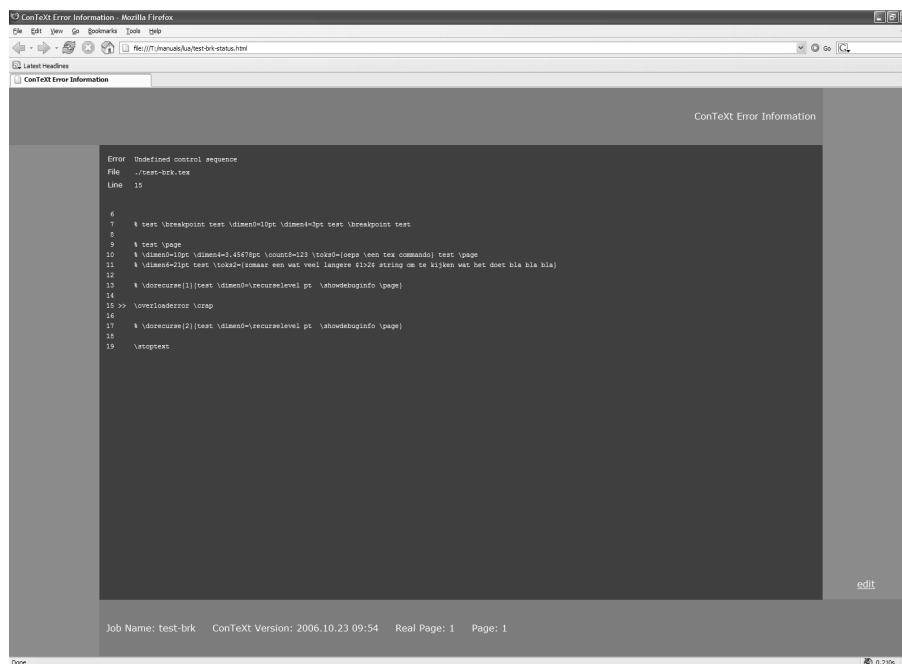


Figure 2. An example error screen.

Playing with such features gives us an impression of what kind of access we need to  $\TeX$ 's internals. It also formed a starting point for conversion routines and a mechanism for embedding LUA code in HTML pages generated by Con $\TeX$ T.

## file io

Replacing  $\TeX$ 's in- and output handling is non-trivial. Not only is the code quite interwoven in the WEB2C source, but there is also the KPSE library to deal with. This means that quite some callbacks are needed to handle the different types of files. Also, there is output to the log and terminal to take care of.

Getting this done took us quite some time and testing and debugging was good for some headaches. The mechanisms changed a few times, and  $\TeX$  and LUA code was thrown away as soon as better solutions came around. Because we were testing on real documents, using a fully loaded Con $\TeX$ T we could converge to a stable version after a while.

Getting this IO stuff done is tightly related to generating the format and starting up L $\TeX$ . If you want to overload the file searching and IO handling, you need overload as soon as possible. Because L $\TeX$  is also supposed to work with the existing KPSE library, we still have that as fallback, but in principle one could think of a KPSE free version, in which case the default file searching is limited to the local path and memory initialization also reverts to the hard coded defaults. A complication is that the source code has KPSE calls and references to KPSE variables all over the place, so occasionally we run into interesting bugs.

Anyhow, while Taco hacked his way around the code, I converted my existing RUBY based KPSE variant into LUA and started working from that point. The advantage of having our own IO handler is that we can go beyond KPSE. For instance,

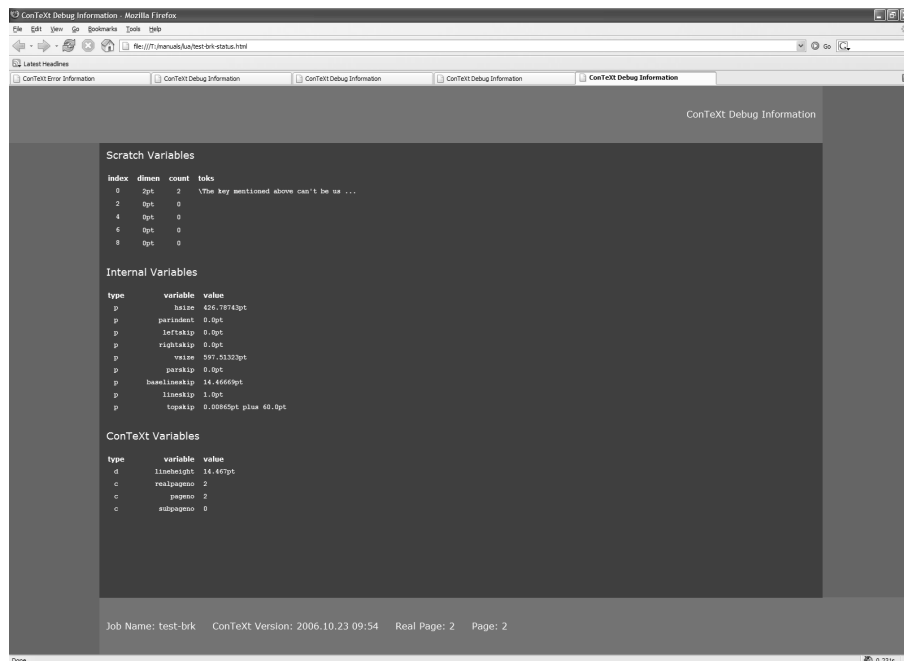


Figure 3. An example debug screen.

since L<sup>A</sup>T<sub>E</sub>X has, among a few others, the ZIP libraries linked in, we can read from ZIP files, and keep all T<sub>E</sub>X related files in TDS compliant ZIP files as well. This means that one can say:

```
\input zip::somezipfile::somefile.tex
\input zip://somezipfile.zip/somepath/somefile.tex
```

and use similar references to access files. Of course we had to make sure that KPSE like searching in the TDS (standardized T<sub>E</sub>X trees) works smoothly. There are plans to link the curl library into L<sup>A</sup>T<sub>E</sub>X, so that we can go beyond this and access repositories.

Of course, in order to be more or less KPSE and WEB2C compliant, we also need to support this paranoid file handling, so we provide mechanisms for that as well. In addition, we provide ways to create sandboxes for system calls.


Getting to intercept all log output (well, most log output) was a problem in itself. For this I used a (preliminary) XML based log format, which will make log parsing easier. Because we have full control over file searching, opening and closing, we can also provide more information about what files are loaded. For instance we can now easily trace what TFM files T<sub>E</sub>X reads.

Implementing additional methods for locating and opening files is not that complex because the library that ships with ConT<sub>E</sub>Xt is already prepared for this. For instance, implementing support for:

```
\input http://www.someplace.org/somepath/somefile.tex
```

involved a few lines of code, most of which deals with caching the files. Because we overload the whole IO handling, this means that the following works ok:

```
\placefigure
{http handling}
{\externalfigure
 [http://www.pragma-ade.com/show-gra.pdf]
 [page=1,width=\textwidth]}
```



**Figure 4.** http handling

Other protocols, like FTP are also supported, so one can say:

```
\typefile {ftp://anonymous:@ctan.org/tex-archive/systems\  
/knuth/lib/plain.tex}
```

On the agenda is playing with database, but by the time that we enter that stage linking the curl libraries into L<sup>A</sup>T<sub>E</sub>X should have taken place.

#### **verbatim**

The advance of L<sup>A</sup>T<sub>E</sub>X also permitted us to play with a long standing wish of catcode tables, a mechanism to quickly switch between different ways of treating input characters. An example of a place where such changes take place is `verbatim` (and in ConT<sub>E</sub>Xt also when dealing with XML input).

We already had encountered the phenomena that when piping back results from LUA to T<sub>E</sub>X, we needed to take care of catcodes so that T<sub>E</sub>X would see the input as we wished. Earlier experiments with applying `\scantokens` to a result and thereby interpreting the result conforming the current catcode regime was not sufficient or at least not handy enough, especially in the perspective of fully expandable LUA results. To be honest, `\scantokens` was rather useless for this purposes due to its pseudo file nature and its end-of-file handling but in L<sup>A</sup>T<sub>E</sub>X we now have a convenient `\scantextokens` which has no side effects.

Once catcode tables were in place, and the relevant ConT<sub>E</sub>Xt code adapted, I could start playing with one of the trickier parts of T<sub>E</sub>X programming: typesetting T<sub>E</sub>X using T<sub>E</sub>X, or `verbatim`. Because in ConT<sub>E</sub>Xt `verbatim` is also related to buffering and pretty printing, all these mechanism were handled at once. It proved to be a pretty good testcase for writing LUA results back to T<sub>E</sub>X, because anything you can imagine can and will interfere (line endings, catcode changes, looking ahead for arguments, etc). This is one of the areas where Mark IV code will make things look more clean and understandable, especially because we could move all kind of postprocessing (needed for pretty printing, i.e. syntax highlighting) to LUA. Interesting is that the resulting code is not beforehand faster.



Pretty printing 1000 small (one line) buffers and 5000 simple `\type` commands perform as follows:

	TeX normal	TeX pretty	Lua normal	Lua pretty
buffer	2.5 (2.35)	4.5 (3.05)	2.2 (1.8)	2.5 (2.0)
inline	7.7 (4.90)	11.5 (7.25)	9.1 (6.3)	10.9 (7.5)

Between braces the runtime on Taco’s more modern machine is shown. It’s not that easy to draw conclusions from this because TeX uses files for buffers and with LUA we store buffers in memory. For inline verbatim, LUA call’s bring some overhead, but with more complex content, this becomes less noticable. Also, the LUA code is probably less optimized than the TeX code, and we don’t know yet what benefits a Just In Time LUA compiler will bring.

### xml

Interesting is that the first experiments with XML processing don’t show the expected gain in speed. This is due to the fact that the ConTeXt XML parser is highly optimized. However, if we want to load a whole XML file, for instance the formal ConTeXt interface specification `cont-en.xml`, then we can bring down loading time (as well as TeX memory usage) down from multiple seconds to a blink of the eyes. Experiments with internal mappings and manipulations demonstrated that we may not so much need an alternative for the current parser, but can add additional, special purpose ones.

We may consider linking `xsltproc` into `LUATeX`, but this is yet undecided. After all, the problem of typesetting does not really change, so we may as well keep the process of manipulating and typesetting separated.

### multipass data

Those who know ConTeXt a bit will know that it may need multiple passes to typeset a document. ConTeXt not only keeps track of index entries, list entries, cross references, but also optimizes some of the output based on information gathered in previous passes. Especially so called two-pass data and positional information puts some demands on memory and runtime. Two-pass data is collapsed in lists because otherwise we would run out of memory (at least this was true years ago when these mechanisms were introduced). Positional information is stored in hashes and has always put a bit of a burden on the size of a so called utility file (ConTeXt stores all information in one auxiliary file).

These two datatypes were the first we moved to a LUA auxiliary file and eventually all information will move there. The advantage is that we can use efficient hashes (without limitations) and only need to run over the file once. And LUA is incredibly fast in loading the tables where we keep track of these things. For instance, a test file storing and reading 10.000 complex positions takes 3.2 seconds runtime with `LUATeX` but 8.7 seconds with traditional `PDFTeX`. Imagine what this will save when dealing with huge files (400 page 300 Meg files) that need three or more passes to be typeset. And, now we can without problems bump position tracking to millions of positions.

### Initialization revised

Initializing `LUATeX` in such a way that it does what you want it to do your way can be tricky. This has to do with the fact that if we want to overload certain features (using callbacks) we need to do that before the originals start doing their work. For instance, if we want to install our own file handling, we must make sure that the built-in file searching does not get initialized. This is particularly important when the built in search engine is based on the `KPSE` library. In that case the first serious file access will result in loading the `ls-R` filename databases, which will take an amount of time more or less linear with the size of the TeX trees. Among the

reasons why we want to replace KPSE are the facts that we want to access ZIP files, do more specific file searches, use HTTP, FTP and whatever comes around, integrate ConT<sub>E</sub>Xt specific methods, etc.

Although modern operating systems will cache files in memory, creating the internal data structures (hashes) from the rather dumb files take some time. On the machine where I was developing the first experimental L<sup>A</sup>T<sub>E</sub>X code, we're talking about 0.3 seconds for PDF<sub>T</sub><sub>E</sub>X. One would expect a L<sup>A</sup> based alternative to be slower, but it is not. This may be due to the different implementation, but for sure the more efficient file cache plays a role as well. So, by completely disabling KPSE, we can have more advanced IO related features (like reading from ZIP files) at about the same speed (or even faster). In due time we will also support progname (and format) specific caches, which speeds up loading. In case one wonders why we bother about a mere few hundreds of milliseconds: imagine frequent runs from an editor or sub-runs during a job. In such situation every speed up matters.

So, back to initialization: how do we initialize L<sup>A</sup>T<sub>E</sub>X. The method described here is developed for ConT<sub>E</sub>Xt but is not limited to this macro package; when one tells T<sub>E</sub>XEXEC to generate formats using the `--luatex` directive, it will generate the ConT<sub>E</sub>Xt formats as well as MPTOPDF using this engine.

For practical reasons, the Lua based IO handler is KPSE compliant. This means that the normal `texmf.cnf` and `ls-R` files can be used. However, their content is converted in a more L<sup>A</sup> friendly way. Although this can be done at runtime, it makes more sense to do this in advance using LUATOOLS. The files involved are:

<b>input</b>	<b>raw input</b>	<b>runtime input</b>	<b>runtime fallback</b>
	<code>ls-R</code>	<code>files.luc</code>	<code>files.lua</code>
<code>texmf.lua</code>	<code>texmf.cnf</code>	<code>configuration.luc</code>	<code>configuration.lua</code>

In due time LUATOOLS will generate the directory listing itself (for this some extra libraries need to be linked in). The configuration file(s) eventually will move to a L<sup>A</sup> table format, and when a `texmf.lua` file is present, that one will be used.

```
luatools --generate
```

This command will generate the relevant databases. Optionally you can provide `--minimize` which will generate a leaner database, which in turn will bring down loading time to (on my machine) about 0.1 sec instead of 0.2 seconds. The `--sort` option will give nicer intermediate (`.lua`) files that are more handy for debugging.

When done, you can use LUATOOLS roughly in the same manner as KPSEWHICH, for instance to locate files:

```
luatools texnansi-lmr10.tfm
luatools --all tufte.tex
```

You can also inspect its internal state, for instance with:

```
luatools --variables --pattern=TEXMF
luatools --expansions --pattern=context
```

This will show you the (expanded) variables from the configuration files. Normally you don't need to go that deep into the belly.

The LUATOOLS script can also generate a format and run L<sup>A</sup>T<sub>E</sub>X. For ConT<sub>E</sub>Xt this is normally done with the T<sub>E</sub>XEXEC wrapper, for instance:

```
texexec --make --all --luatex
```

When dealing with this process we need to keep several things in mind:

- L<sup>A</sup>T<sub>E</sub>X needs a L<sup>A</sup> startup file in both ini and runtime mode
- these files may be the same but may also be different

- here we use the same files but a compiled one in runtime mode
- we cannot yet use a file location mechanism

A `.luc` file is a precompiled LUA chunk. In order to guard consistency between LUA code and tex code, ConT<sub>E</sub>Xt will preload all LUA code and store them in the bytecode table provided by L<sub>U</sub>A<sub>T</sub>E<sub>X</sub>. How this is done, is another story. Contrary to these tables, the initialization code can not be put into the format, if only because at that stage we still need to set up memory and other parameters.

In our case, especially because we want to overload the IO handler, we want to store the startup file in the same path as the format file. This means that scripts that deal with format generation also need to take care of (relocating) the startup file. Normally we will use T<sub>E</sub>X<sub>E</sub>EXEC but we can also use LUATOOLS.

Say that we want to make a plain format. We can call LUATOOLS as follows:

```
luatools --ini plain
```

This will give us (in the current path):

```
120,808 plain.fmt
 2,650 plain.log
 80,767 plain.lua
 64,807 plain.luc
```

From now on, only the `plain.fmt` and `plain.luc` file are important. Processing a file

```
test \end
```

can be done with:

```
luatools --fmt=./plain.fmt test
```

This returns:

```
This is luaTeX, Version 3.141592-0.1-alpha-20061018 (Web2C 7.5.5)
(./test.tex [1] )
Output written on test.dvi (1 page, 260 bytes).
Transcript written on test.log.
```

which looks rather familiar. Keep in mind that at this stage we still run good old Plain T<sub>E</sub>X. In due time we will provide a few files that will making work with LUA more convenient in Plain T<sub>E</sub>X, but at this moment you can already use for instance `\directlua`.

In case you wonder how this is related to ConT<sub>E</sub>Xt, well only to the extend that it uses a couple of rather generic ConT<sub>E</sub>Xt related LUA files.

ConT<sub>E</sub>Xt users can best use T<sub>E</sub>X<sub>E</sub>EXEC which will relocate the format related files to the regular engine path. In LUATOOLS terms we have two choices:

```
luatools --ini cont-en
luatools --ini --compile cont-en
```

The difference is that in the first case `context.lua` is used as startup file. This LUA file creates the `cont-en.luc` runtime file. In the second call LUATOOLS will create a `cont-en.lua` file and compile that one. An even more specific call would be:

```
luatools --ini --compile --luafile=blabla.lua          cont-en
luatools --ini --compile --lualibs=bla-1.lua,bla-2.lua cont-en
```

This call does not make much sense for ConT<sub>E</sub>Xt. Keep in mind that LUATOOLS does not set up user specific configurations, for instance the `--all` switch in T<sub>E</sub>X<sub>E</sub>EXEC will set up all patterns.

I know that it sounds a bit messy, but till we have a more clear picture of where L<sup>A</sup>T<sub>E</sub>X is heading this is the way to proceed. The average ConT<sub>E</sub>Xt user won't notice those details, because T<sub>E</sub>XEXEC will take care of things.

Currently we follow the TDS and WEB2C conventions, but in the future we may follow different or additional approaches. This may as well be driven by more complex IO models. For the moment extensions still fit in. For instance, in order to support access to remote resources and related caching, we have added to the configuration file the variable:

```
TEXMFCACHE = $TMP;$TEMP;$TMPDIR;$HOME;$TEXMFVAR;$VARTEXMF; .
```

## An example: CalcMath

### introduction

For a long time T<sub>E</sub>X's way of coding math has dominated the typesetting world. However, this kind of coding is not that well suited for non academics, like schoolkids. Often kids do know how to key in math because they use advanced calculators. So, when a couple of years ago we were implementing a workflow where kids could fill in their math workbooks (with exercises) on-line, it made sense to support so called Texas Instruments math input. Because we had to parse the form data anyway, we could use a `[ [ and ] ]` as math delimiters instead of `$`. The conversion too place right after the form was received by the web server.

<code>sin(x) + x^2 + x^(1+x) + 1/x^2</code>	$\sin(x) + x^2 + x^{1+x} + \frac{1}{x^2}$
<code>mean(x+mean(y))</code>	$\overline{x + \bar{y}}$
<code>int(a,b,c)</code>	$\int_b^a c$
<code>(1+x)/(1+x) + (1+x)/(1+(1+x)/(1+x))</code>	$\frac{1+x}{1+x} + \frac{1+x}{1+\frac{1+x}{1+x}}$
<code>10E-2</code>	$10 \times 10^{-2}$
<code>(1+x)/x</code>	$\frac{1+x}{x}$
<code>(1+x)/12</code>	$\frac{1+x}{12}$
<code>(1+x)/-12</code>	$\frac{1+x}{-12}$
<code>1/-12</code>	$\frac{1}{-12}$
<code>12x/(1+x)</code>	$\frac{12x}{1+x}$
<code>exp(x+exp(x+1))</code>	$e^{x+e^{x+1}}$
<code>abs(x+abs(x+1)) + pi + inf</code>	$ x +  x + 1   + \pi + \text{inf}$
<code>Dx Dy</code>	$\frac{dx}{dx} \frac{dy}{dx}$
<code>D(x+D(y))</code>	$\frac{d}{dx}(x + \frac{d}{dx}(y))$
<code>Df(x)</code>	$f'(x)$
<code>g(x)</code>	$g(x)$
<code>sqrt(sin^2(x)+cos^2(x))</code>	$\sqrt{\sin^2(x) + \cos^2(x)}$

By combining L<sup>A</sup> with T<sub>E</sub>X, we can do the conversion from calculator math to T<sub>E</sub>X immediately, without auxiliary programs or complex parsing using T<sub>E</sub>X macros.

**tex**

In a ConT<sub>E</sub>Xt source one can use the `\calcmath` command, as in:

The strange formula `\calcmath {\sqrt{\sin^2(x)+\cos^2(x)}}` boils down to ...

One needs to load the module first, using:

```
\usemodule[calcmath]
```

Because the amount of code involved is rather small, eventually we may decide to add this support to the Mark IV kernel.

**xml**

Coding math in T<sub>E</sub>X is rather efficient. In XML one needs way more code. Presentation MATHML provides a few basic constructs and boils down to combining those building blocks. Content MATHML is better, especially from the perspective of applications that need to do interpret the formulas. It permits for instance the ConT<sub>E</sub>Xt content MATHML handler to adapt the rendering to cultural driven needs. The OPENMATH way of coding is like content MATHML, but more verbose with less tags. Calculator math is more restrictive than T<sub>E</sub>X math and less verbose than any of the XML variants. It looks like:

```
<icm>\sqrt{\sin^2(x)+\cos^2(x)}</icm> test
```

And in display mode:

```
<dcM>\sqrt{\sin^2(x)+\cos^2(x)}</dcM> test
```

**speed**

This script (which you can find in the ConT<sub>E</sub>Xt distribution as soon as the Mark IV code variants are added) is the first real T<sub>E</sub>X related LUA code that I wrote; so far I had only written some wrapping and spell checking code for the SciTE editor. It also made a nice demo for a couple of talks that I held at usergroup meetings. The script has a lot of expressions. These convert one string into another. They are less powerful than regular expressions, but pretty fast and adequate. The feature I miss most is alternation like `(1|st)uck` but it's a small price to pay. As the LUA manual explains: adding a POSIX compliant regex parser would take more lines of code than LUA currently does.

On my machine, running this first version took 3.5 seconds for 2500 times typesetting the previously shown square root of sine and cosine. Of this, 2.1 seconds were spent on typesetting and 1.4 seconds on converting. After optimizing the code, 0.8 seconds were used for conversion. A stand alone LUA takes .65 seconds, which includes loading the interpreter. On a test of 25.000 sample conversions, we could gain some 20% conversion time using the LUAJIT just in time compiler.

Hans Hagen