

T_EX forever!

Jonathan Fine
 Learning and Teaching Solutions
 The Open University
 Milton Keynes
 United Kingdom
 J.Fine@open.ac.uk
<http://www.pytex.org>

Abstract

This paper explores new ways of doing input to and output from T_EX. These new ways bypass our current habits, and provide fresh opportunities.

Usually, T_EX is run as a batch program. But when run as a daemon, T_EX can be part of an interactive program. Daemons often that run forever, or at least for a long time. Hence the title of this paper.

Usually, parsing and transformation of the input data is done by T_EX macros, with little outside help. Often, this results in input documents that only T_EX can understand. Also, T_EX macros can be hard to write. We demonstrate the replacement of T_EX macros by an external program. This is done in real time.

Usually, T_EX's principal output is a dvi representation of typeset pages, for processing by a printer driver. However, T_EX's log file or console can be used to allow T_EX to output the boxes it holds internally. (Alternatively, an extension of T_EX could write this data out in a binary form.) Shipping out boxes rather than dvi allows an external program to do the page makeup.

Don Knuth's original conception was that T_EX would be "just a typesetting language". In some sense he "put in many of T_EX's programming features only after kicking and screaming". The developments described above reduce our dependence on T_EX macros, and so bring our use of T_EX closer to Knuth's original conception. Doing this will greatly improve its usefulness.

Long live T_EX!

Introduction

In 1990, Don Knuth told us [8, p.572] that his work on developing T_EX had come to an end. He went on to say:

Of course I do not claim to have found the best solution to every problem. I simply claim that it is a great advantage to have a fixed point as a building block. Improved macro packages can be added on the input side; improved device drivers can be added on the output side.

In this paper, the author tries to follow this advice. There are imperfections in T_EX, and the lack of proper support for Unicode fonts and filenames is a major weakness. However, T_EX also has enormous strengths. It is archival. It carefully uses integer arithmetic to ensure that it gets the same line and page breaks, regardless of the machine it is running

on. Its algorithm for breaking a paragraph into lines is reliable, adaptable and efficient. T_EX is without rival for complex mathematical typesetting.

Often, T_EX is used with L^AT_EX as the macro package front end, and with dvips as the device driver. Sometimes, the word 'T_EX' is used to refer to the whole system. However, in this paper we mean by 'T_EX' the typesetting program written by Don Knuth. And so L^AT_EX and dvips are tools for use with T_EX.

This paper is concerned with making improvements on the input and output sides of T_EX, both areas of work where there is an enormous amount still to do. However, our proposals are not exactly macro packages and device drivers.

A note to the reader: This paper has been written for a general audience, and in particular for those who are not T_EX experts. At the same time, discussion of technical details is at times either unavoidable

1	¶	textfile → marked up text and math
2	†	text + transform → horizontal primitives
3	*	horizontal primitives → hlist
4	†	math + transform → math primitives
5	*	math primitives + parameters → hlist
6	*	hlist + parameters → vlist
7	†	vlists + page make up → page boxes
8	*	page boxes → sequential dvi file
9	¶	sequential dvi file → random access dvi file
10	¶	random access dvi file → rendered page

Table 1: How T_EX works, in 10 stages

† usually done using T_EX macros.

* usually done using T_EX’s built in procedures.

¶ file input and output matters.

able or helpful. Therefore, I hope that the experts will forgive my stating the obvious, and that the others forgive my discussing the difficult.

References: Many of the articles cited here have been reprinted in the collection *Digital Typography* [13]. Page numbers in citations refer to [13], and not to the original publication.

How T_EX works

Table 1 gives a concise description as to how T_EX works. On the input side we propose that an external program perform the transformation in steps 2 and 4.

On the output side we have two proposals. The first is that (8) be replaced by:

8’. page boxes → stream of dvi pages

The second, which is more ambitious, is that (7) be replaced by:

7’. vlist → external program

followed by page makeup in that external program.

Thus we continue to use T_EX’s excellent typesetting, but reduce the use of its macros.

T_EX — just a typesetting language

T_EX is a typesetting program, written by Don Knuth, that is particular good at mathematical and technical typesetting. T_EX is reliable and stable, and is very widely used by academic mathematicians and physicists.

T_EX has a macro programming language, which allows features to be added. The best known and most widely used T_EX macro package is L^AT_EX. (This is not quite accurate. Although originally L^AT_EX used T_EX, since 2003 it by default uses e-T_EX, which is an extension of T_EX. So it is no longer purely a T_EX macro package. This has no bearing on our discussion.)

In 1996 Don Knuth, describing his intentions when he started to develop T_EX, said [11, p.648]:

I’m not going to design a programming language; I want to have just a typesetting language.

and at the same time he said (loc. cit.):

In some sense I put in many of T_EX’s programming features only after kicking and screaming. [...] In the 70s, I had a negative reaction to software that tries to be all things to all people. Every system had its own universal Turing machine built into it somehow, and everybody’s machine was a little different from everybody else’s.

But the need for more features caused the programming constructs to grow (see Table 2 below). See also [16] for a ‘wish-list’ of future developments.

Therefore, by *removing* commands from T_EX, we can come closer to Don’s original conception of T_EX. However, for this to succeed in practice, some other means of adding new features is required. Indeed, one of the major problems T_EX users have now is that the existing programming constructs barely support the demand for new features. This we discuss later.

In this section we outline how to cut T_EX down to the bare minimum. To be specific, in this section we ask: What commands are required in order to access T_EX’s algorithm for breaking a paragraph into lines?

To create a paragraph one needs to be able to load fonts, change fonts, and set a character in the current font. One also needs commands for appending glue, kerns and the like to the paragraph.

To break the paragraph into lines, one needs the `\par` primitive (also known as `\endgraf`) and a means of assigning values to the line-breaking parameters, such as `\hsize`.

In other words, the basic operations are to add an item to the current horizontal list, and to form a paragraph out of the current horizontal list. (For mathematics and table typesetting there are similar basic operations.)

It should at this point be clear that certain primitive T_EX commands are *not* required in order to do typesetting. These commands include all the `\def` commands (such as `\def`, `\chardef`, `\xdef`), `\let`, `\begingroup` and `\endgroup`. Once category codes have been set up, there is no further need for `\catcode`. And there is certainly no need for commands such as `\expandafter`, `\noexpand`, `\aftergroup` and `\futurelet`. All these are *not*

Control sequence	Date added
<code>\if</code>	21 Jun 1978
<code>\pausing</code>	16 Mar 1978
<code>\uppercase</code>	25 Nov 1978
<code>\xdef</code>	28 Nov 1978
<code>\ifmode</code>	23 July 1978
<i>active characters</i>	25 Jan 1980
<code>\let</code>	25 Mar 1980
<code>\ifx</code>	13 July 1981
<code>\catcode</code>	July 1982
<code>\expandafter, \openin</code>	12 Sep 1982
<code>\string</code>	12 Sep 1982
<code>\immediate</code>	12 Oct 1982
<code>\csname, \endcsname, \fi</code>	13 Nov 1982
<code>\everymath, \everydisplay, \futurelet</code>	2 Dec 1982
<code>\endinput</code>	7 Dec 1982
<code>\jobname</code>	25 Dec 1982
<code>\globaldefs</code>	20 Jan 1983
<code>\iffalse, \iftrue</code>	3 Feb 1983
<code>\everyvbox, \everyhbox</code>	6 Mar 1983
<code>\everyjob</code>	18 Mar 1983
<code>\advance, \multiply, \divide \noexpand, \meaning</code>	25 May 1983
<code>\afterassignment</code>	27 May 1983
<code>\escapechar, \endlinechar</code>	4 Jul 1983
<code>\errhelp</code>	11 Jul 1983
<code>\aftergroup, \newlinechar</code>	16 Jul 1983
<code>\ifhbox, \ifvbox</code>	27 Aug 1983
<code>\holdinginserts</code>	30 Sep 1989

Table 2: Some T_EX control sequences not needed for typesetting (after [7])

typesetting commands, and exist only to allow features to be added to T_EX.

Moreover, these commands cause difficulty for both T_EX users and programmers. Their introduction is perhaps a sign that things were starting to go in the wrong direction.

In [7] Knuth published, in edited form, the log books he kept while he was developing T_EX. In these, we can see the introduction of features. (See Table 2.)

Suppose all programming commands are disabled by `\let`-ting them be undefined, like so:

```
\let \afterassignment = \undefined
```

Provided we remove enough commands, we will have, as Don wanted T_EX to be in the first place, “just a typesetting language”. A language without features, and without the capability of adding features (which is itself a feature).

Comparison with PostScript and with machine code is instructive. Most PostScript is generated by programs that translate from a higher-level language

down to PostScript. Similarly, much machine code is generated by compiling ‘C’ source files.

Many of us write input files for (L^A)T_EX, using a text editor. We won’t do that for a featureless T_EX. It’s too much hard work, and anyway we want to write in a higher-level language. We are suggesting that an external program perform the text transformation that is traditionally performed by a T_EX macro package.

Improved macros — input transformation

This section could also be titled:

```
\let \def = \undefined
```

Don suggested that we add improved macro packages on the input side. Now, a macro package has two main purposes. One is issuing typesetting instructions to T_EX. This will create a galley (or page of unlimited depth). The second purpose is the output (or page makeup) routine, which breaks the galley into pages of a suitable size.

In this section we consider the creation of a galley. Marked-up text, such as

```
\section{Improved macros}
```

is translated (by L^AT_EX in this case) into a large number of low-level instructions. The title

```
Improved macros
```

is scarcely translated. Each character sets itself, and space characters produce default interword glue. (Later, we present an example of this.)

It is `\section{}` that does most of the work. Here are some of the technical details. It selects the font to be used, and the paragraph parameters for the title (in case it is wider than the measure). It also places glue and penalties before and after the title on the galley. It might also add a section number, and record information for the table of contents.

High-level commands are being translated into low-level typesetting instructions. This translation need not be done by L^AT_EX (or indeed by any other T_EX macro package). For example, in the WEB system of literate programming, much of the work is done using external programs. Similarly, XSLT templates are often used to transform text, prior to it being passed to T_EX to typesetting.

For over 10 years the L^AT_EX3 project has been working to enhance L^AT_EX by providing [15, p.1]

a flexible interface for typographic designers to easily specify the formatting of a class of documents.

Such an interface might, for example, be similar to Cascading Style Sheets (CSS) for HTML. We have seen that Don Knuth only reluctantly added

programming features to T_EX. The author believes that T_EX macros are not a suitable language for creating the above interface, and that the long delay in its delivery is evidence for this.

This interface could instead be written as an external program. In a later section we describe Q_AT_EX, which is a wrapper around T_EX that effectively allows T_EX to interact with external programs.

Improved macros — output routines

This section could also be titled:

```
\let \output = \undefined
```

One of the most interesting and best parts of T_EX is the algorithm it uses for breaking a paragraph into lines. The algorithm for breaking the galley into pages is not so good, although for simple technical material it is more than adequate.

In this paper we do not suggest improved output algorithms (we have discussed this elsewhere [4]). Instead, we describe a solution to a related problem. The T_EX macro language is not a suitable environment for the writing of complicated page makeup algorithms. Here we describe a means of moving the problem to another domain.

The galley produced by T_EX consists of a vertical list. This vertical list consists of boxes, glue, penalties, and so forth (see *The T_EXbook*, page 110 for a complete list). The `\showlists` command prints a detailed description of the content of this vertical list.

The output of `\showlists` can be parsed by an external program, and used to reconstruct within that program the vertical list created by T_EX. If the external program can also send low-level typesetting instructions to T_EX, then T_EX in effect has become a callable function available to the external program. (As in Q_AT_EX, the interactive console or more exactly stdin and stdout can be used for this communication.)

This is not a completely new idea. In 1996, Jiří Veselý asked [10, pp620–621]:

Once I was asked about the possibility to make a list of all hyphenated words in the book. I was not able to find a way in your book to do this.

To this, Don replied (loc. cit.):

This would be easy to do now in a module specially written for T_EX. I would say that right now, in fact, you could get almost exactly what you want by writing a filter that says to T_EX “Turn on all the tracing options that cause it to list the page con-

tents.” Then a little filter program would take the trace information through a UNIX pipe and it would give you the hyphenated words. It would take an afternoon to write this program; well, maybe two afternoons . . . and a morning.

We develop this idea later in the paper.

Instant Preview and T_EX as a daemon

This section could also be titled:

```
\let \end = \undefined
```

Interactive programs typically require a response time of less than a tenth of a second, while a response in a hundredth of a second is seen as instantaneous.

On my current 800 MHz PC, the command

```
$ tex story \end
```

takes about 0.137 seconds, while

```
$ tex \end
```

takes 0.133 seconds. The first command typesets a small page of material; the second does nothing but start T_EX and then exit. Thus, typesetting the small page takes about 0.004 seconds.

It follows from this that typesetting material for Instant Preview is tolerable if the start-up time is included, while it will be perceived as instantaneous if T_EX is run as a daemon.

Running T_EX as a daemon is an example of T_EX forever. We wish for T_EX to start up when the computer boots, and to remain running indefinitely. Moreover, we might prefer that there were only a single instance of the T_EX daemon running.

Documents and macro packages may have to be adapted, to make the most of Instant Preview. The key concept seem to be this: That the source file be partitioned into regions by markers, which we call ‘belays’, and that the macro package be able to typeset each region independently. In other words, that the macro package support random access typesetting. This is, again, an example of T_EX being enhanced by an improved macro package on the front end.

The author has already written [5] about Instant Preview. At the conference he hopes to demonstrate the latest progress.

Decorating dvi files

T_EX has no built-in notion of colour, or of graphics inclusion. However, the `\special` command allows device drivers to produce special effects. By decoration we mean the application of colour, change bars and the like to the rendered page.

In the domestic setting, decoration of a room or a house does not move the walls or make other structural alterations. In typesetting, adding decoration should not affect typesetting decisions, such as the line breaks and the placement of items on the page. (This is not to say that the typographic design should not take into account the subsequent application of decoration.)

The current practice regarding decoration is to use a fairly simple dvi processor, and to have the (L^A)T_EX macro package place appropriate `\special` commands into the dvi file. From the point of view of a device-driving dvi processor, this is probably correct. It seems that, historically, low-level capabilities were added to device drivers. Then macro packages were written to access these new features.

From the point of view of the macro package, this approach is probably wrong. As already noted, decoration should not affect typesetting decisions. This is an important property, whose fulfilment should be central to the approach taken.

Suppose, for example, that some text is to be printed in a spot colour. Placing a `\special` at the start or end of a word does not affect its hyphenation. Therefore something like

```
% usage: \color{red}{Text to go in red}
\def\color#1#2{%
  \special{push #1}%
  #2%
  \special{pop}%
}
```

will suffice, at least in the simplest cases.

However, a page boundary might break the red text. This places a burden on the dvi processor, to keep track of this information. Typical dvi processors allow random access to the pages in the dvi file. Having to look at previous page(s) breaks this random access.

The solution we suggest is to write a dvi-to-dvi filter that resolves these random access problems. Such a filter is not, of course, a device driver, but it can be used with any device driver. Its purpose is to translate high-level specials into low-level specials.

T_EX is being held back by the weakness of tools for decorating text. For example, a common requirement is to place a background rectangle behind a paragraph of text. If the paragraph is broken over two pages, the background rectangle should be similarly broken.

In 1987 Don Knuth and Pierre MacKay discussed a similar problem, namely implementing bi-directional typesetting without extending T_EX. They wrote [14, p159]:

How can we get T_EX to do this? The best approach is probably to extend the driver programs that produce printed output from the dvi files that T_EX writes, instead of trying to do tricky things with T_EX macros.

In the same article [13, p161] they then produced a “much more reliable and robust scheme by building a specially extended version of T_EX”.

Q_AT_EX — or ask a friend a question

T_EX is a typesetting language, with limited text-processing and other capabilities. Things that are easy to do in other languages are hard to do in T_EX. Examples are to find the dimensions of an EPS or other graphics file, or to calculate the sine and cosine of an angle, so that space can be left for rotated text.

Traditionally, such questions have been answered by writing T_EX macros. The author finds that T_EX macros are not a suitable language for such text manipulation tasks.

Here is an extract from the `\Gread@eps` macro in the L^AT_EX file `graphics.sty`.

```
\immediate\openin\@inputcheck#1 %
\ifeof\@inputcheck
  \@latex@error{File ‘#1’ not found}%
  \@ehc
\else
  \Gread@true
  \let\@tempb\Gread@false
  \loop
    \read\@inputcheck to\@tempa
    \ifeof\@inputcheck
      \Gread@false
    \else
      \expandafter\Gread@find@bb
      \@tempa:.\%
    \fi
  \ifGread@
  \repeat
\immediate\closein\@inputcheck
\fi
```

The author has developed Q_AT_EX, which allows the T_EX macro programmer to ask another process to answer questions. Such as: What is the bounding box of `myfile.eps`?

Q_AT_EX (pronounced ‘kwa-tech’) provides a different route. Questions and answers are the essence of Q_AT_EX. When Q_AT_EX is used, lines such as:

```
!Q=qatexlib.eps.bbox(myfile.eps)
!A=0,0,0 0 35 97
```

appear in the T_EX’s log file.

The first line is a question, produced using a `\write` command. The second line is the answer.

The characters `!A=` are a prompt, produced using a `\message` command.

The remainder of that line is the answer to the question. The prefix `0,0,` tells T_EX that the question was successfully posed and answered. There follows the information asked for. This information is supplied by a process external to T_EX.

T_EX uses `\read -1 to \temp` to read this information from its stdin stream. So far as T_EX is concerned, this data might have come from the user. In fact, it has come from a program, namely Q_AT_EX.

Q_AT_EX works as follows. It is a wrapper program that invokes T_EX, and takes control of its standard input and output. When it sees a question line, followed by the answer prompt, it parses and answers the question, then sends the answer to T_EX, using T_EX's standard input. However, Q_AT_EX does not answer the question itself. It imports a module—in the example above the `eps` module—to answer the question for it.

Here is the definition, in Python, of a function that returns the bounding box of an EPS file, as a string. If not found, it raises an exception. It took me less than 10 minutes to write. It would be part of a `eps` module for use with Q_AT_EX.

```
# File: qatexlib/eps.py
_bb_prefix = '%BoundingBox: '
def qa_bbox(filename):
    f = open(filename)
    for line in f:
        if line.startswith(_bb_prefix):
            sizes = line.split()[1:]
            return ' '.join(sizes)
    msg = "File '%s' has no bounding box"
    raise Exception, msg % filename
```

Alternatives and complements to Q_AT_EX

In this section we compare Q_AT_EX to shell escape, `eval4tex`, PerlT_EX, `sTEXme` and Pymacs. Each of these has some similarity with Q_AT_EX.

Shell escape. Modern implementations allow T_EX to issue shell commands, as if they had been typed at a command prompt. This allows, for example, a command such as `makeindex` to be run after T_EX has processed the body of a document, but before setting the back matter.

However, it also allows other commands to be run, such as the deletion of files. And T_EX documents can execute arbitrary T_EX commands. (Strictly speaking, this is not true. For example, in Active T_EX [3] all characters are active. This allows a macro package to prevent execution directly from a document of all but specified commands.)

Often, T_EX documents are distributed in source form. If shell escape is enabled, the typesetting a document could result in a shell escape command being run, that finds and deletes all your files. Clearly, shell escape is a security risk. For this reason, shell escape is disabled by default, and is rarely used.

Even without shell escape, T_EX macros and therefore documents can overwrite existing files. (T_EX has no inbuilt ability to delete files. But it can destroy their contents.) Therefore, modern implementations of T_EX refuse to open for writing files that are not in or beneath the current directory, or a similarly specified area.

This restriction is not applied to the reading of files. Therefore, it is possible for a T_EX document, when typeset, to include in it other files. These other files might be confidential.

Therefore, in line with the theme of T_EX being “just a typesetting system”, it might be sensible to:

```
\let \openout = \undefined
\let \openin = \undefined
\let \input = \undefined
```

and have another program take responsibility for security. The security monitor could then send material to be typeset to T_EX's standard input stream.

eval4T_EX (by Dorai Sitaram) is a two-pass process that allows T_EX to send expressions to Scheme for evaluation [1]. It provides a `\eval` macro, that is used as below. (The example is Sitaram's, and my exposition follows his).

```
\eval{(display (acos -1))} % gives pi
```

On the first pass, the Scheme code

```
(display (acos -1))
```

is written to an auxiliary file, together with some indexing information.

Before the second pass, a helper program `eval4tex` runs Scheme on the auxiliary file, to create a second auxiliary file. On the second pass, the `\eval` command picks up values from the second auxiliary file, and refreshes the data in the first.

As Sitaram writes:

This strategy is quite common in the T_EX world. The popular T_EX-support programs BibT_EX and MakeIndex, which generate bibliographies and indices respectively, both operate this way.

sT_EXme (by Oleg Paraschenko) is another approach to integrating T_EX with Scheme [19]. Here is his summary of the goals of the project.

The (L^A)T_EX macro language was a great development when it appeared, but now it is

too out-of-date. Programming in T_EX is a fun, but more often it is a pain.

As it seems for me, only very few people can write (L^A)T_EX macros, but a lot of people would like doing it (like me, for example). This is the problem.

One of the solutions is to provide another scripting language for T_EX. That's what is the goal of the sT_EXme project. It should provide the Scheme programming language as a T_EX scripting language.

This project has two parts, namely an extension to T_EX, that allows it to interpret Scheme code, and an extension to Scheme that allow it access T_EX internals. We say more on this later.

PerlT_EX (by Scott Pakin) uses standard Perl and T_EX without extensions [17]. Here is his summary of the goals of the project.

T_EX is a professional-quality typesetting system. However, its programming language is rather hard to use for anything but the most simple forms of text substitution. Even L^AT_EX, the most popular macro package for T_EX, does little to simplify T_EX programming.

Perl is a general-purpose programming language whose forte is in text manipulation. However, it has no support whatsoever for typesetting.

PerlT_EX's goal is to bridge these two worlds. It enables the construction of documents that are primarily L^AT_EX-based but contain a modicum of Perl. PerlT_EX seamlessly integrates Perl code into a L^AT_EX document, enabling the user to define macros whose bodies consist of Perl code instead of T_EX and L^AT_EX code.

Here is Scott Pakin's equivalent to `\eval`:

```
\perlnewcommand{\reversewords}[1]
  {join " ", reverse split " ", $_[0]}
\reversewords{TeX forever!}
```

PerlT_EX, like Q_AT_EX, invokes T_EX under the control of a separate process. Unlike Q_AT_EX, it does not take control of T_EX's standard input and output. Invoking `\reversewords` causes T_EX to write material to a specially named file. This file corresponds to the question in Q_AT_EX. The controlling Perl process then computes the answer to the question, and writes it to another specially named file. Meanwhile, the T_EX process goes into a loop, to poll for the existence of the answer file. Once found, it is `\input` by T_EX.

PerlT_EX seems to have a performance problem. On my 800 Mhz PC, the following example:

```
\documentclass{article}
\usepackage{perltex}
\perlnewcommand{\nothing}{}

\begin{document}
% I've got plenty of nothing ...
\nothing\nothing\nothing
\nothing\nothing\nothing
\nothing\nothing\nothing
\nothing          % 10 nothings
% We're busy doing nothing ...
\end{document}
```

takes about 3.0 seconds to execute. This includes startup time. (On the same machine, it takes Q_AT_EX about 1/3000 seconds to do nothing once.)

Here is at least a partial explanation. Instrumenting the code for PerlT_EX shows that in compiling the above document, T_EX polls for the existence of the helper file approximately 5,000 times. The exact number varies. Adding:

```
\input nothing % input an empty file
```

to the polling loop *reduces* the time taken to about 1.8 seconds, and reduces the number of pollings to about 500. The UNIX `nice` command could also help here.

Pymacs (by François Pinard) is not a way of using Python with T_EX. It is a way of using Python with the Emacs editor [18]. To quote its author:

Pymacs is a powerful tool which, once started from Emacs, allows both-way communication between Emacs Lisp and Python. Pymacs aims Python as an extension language for Emacs rather than the other way around, and this asymmetry is reflected in some design choices. Within Emacs Lisp code, one may load and use Python modules. Python functions may themselves use Emacs services, and handle Emacs Lisp objects kept in Emacs Lisp space.

Pymacs is mentioned because it was highly influential on the author's approach to the integration of T_EX with a scripting language. (At that time, Python had not been chosen.)

Different approaches compared

In the previous two sections we looked at Q_AT_EX, shell escape, `eval4tex`, sT_EXme and PerlT_EX. In this section we make some comparisons.

Philosophy Q_AT_EX is like shell escape, in that simple queries are sent to another process. The other approaches assume that complex code will be written within, or otherwise produced by, T_EX macros. This code is then evaluated by another program.

For example, in Q_AT_EX the problem of reversing words might result in the following conversation between T_EX and the external process:

```
!Q=qatexlib.string.reverse(TeX forever!)
!A=0,0,forever! TeX
```

The question sent to Q_AT_EX could not be along the lines of: “What is the result of evaluating this complex Perl or Scheme expression?” However, such is not the expected use. Rather, it is expected that T_EX will send a short and simple query. If the answer is long, it could be placed in an external file. Once that is done, T_EX can be told, as the answer, that the file is ready to be `\input` (assuming `\input` is still defined).

Architecture The architecture of the implementation of PerlT_EX is closer to that of `eval4tex` than that of `sTeXme`. Perl code is placed in the body of T_EX macros, and this code is sent out to Perl for evaluation. Unlike `sTeXme`, and like `eval4tex`, it does not require either an extension to T_EX or a special version of the command interpreter for the extension language.

PerlT_EX is similar to Q_AT_EX in that T_EX is run under the control of an external program. However, Q_AT_EX uses standard input and output for communication, whereas PerlT_EX polls named files.

Q_AT_EX provides an efficient and portable low-level interface between T_EX and an external process. PerlT_EX is a higher-level package. There is no reason why the Q_AT_EX interface, or something similar to it, should not be used by PerlT_EX, so as to improve performance. The same applies to `eval4tex`, where using this interface would remove the need for a second run. It would also provide better interaction.

Security Any process that evaluates code supplied by a document will expose a security problem, unless the code evaluator is already secure (as in Java, for example). PerlT_EX provides security by using a secure Perl sandbox. Q_AT_EX provides security by having T_EX (and thus the document) supply only data.

Of course, any program that evaluates untrusted data as if it were trusted code has a security issue. If it is necessary to evaluate safely untrusted code, then a secure sandbox is required. This is the key security issue. Q_AT_EX is a small interface

application, which does not attempt to solve this problem. There is no reason why Q_AT_EX should not be used with such a secure sandbox. But that is a matter for the developer who builds upon Q_AT_EX.

Integration and extension Of all the projects considered in this section, `sTeXme` is the most ambitious. It involves making major extensions to T_EX, to produce a new program, called `sTeXme`.

The new name is necessary. T_EX experts will already know that Don Knuth does not object to the sources of T_EX the program being used as the basis for creating a superior program. However, he is most insistent that programs that are not T_EX must not be called T_EX. More exactly, in [8, p572] he wrote:

That is all I ask, after devoting a substantial portion of my life to the creation of these systems and making them available to everybody in the world. I sincerely hope that the members of TUG will help me to enforce these wishes, by putting severe pressure on any person or group who produces any incompatible system and calls it T_EX or METAFONT or Computer Modern — no matter how slight the incompatibility might seem.

This insistence on the stability and consistency of T_EX is, in this author’s view, a significant contribution to its longevity. Users know what to expect, and get what they expect, when they use T_EX.

Regarding the scope of his new program `sTeXme`, Oleg Paraschenko reports:

[...] Scheme code can be executed from TeX and that Scheme code can access TeX internals (getting a string from the TeX string pool, getting a macro definition as the Scheme list).

The source file `stexmelib.c` on the SourceForge repository indicates that Paraschenko is building a C-coded Scheme extension, in which equivalents to T_EX boxes and the like can be stored. This indicates that there are many points of contact between his project and the next section.

PyT_EX

The goal of the PyT_EX project is to combine Python programming with T_EX typesetting. To understand this, think of T_EL/T_EK: T_EL is a scripting language and Tk is a toolkit for building GUI programs. Perl and Python also have interfaces to Tk, allowing them to use Tk when building GUI programs.

Now think of L^AT_EX as L^A/T_EX. L^A is a front end to the T_EX typesetting program written in T_EX’s

macro language. PyT_EX, or Py/T_EX if you prefer, is to be a front end to T_EX written in Python.

PyT_EX replaces the part of T_EX that Don Knuth said he did not want to write, namely the T_EX macro programming language, with something more widely used. Our aim is to provide an alternative means of programming typesetting decisions and logic. This will make T_EX easier to use.

Here is an example of the interface. We are in Python, and wish to typeset a paragraph of text, namely

The cat sat on the mat.

to a measure of 6 picas, which is about one inch (or 25 millimetres). This is a toy example. After typesetting the paragraph, we wish to bring it into Python for page makeup.

To see how this is done, issue the command

```
$ cat cat_sat.tex | tex | grep '^[\.]'
```

where the source stream is

```
% cat_sat.tex
\tracingonline 1
\showboxbreadth \maxdimen
\showboxdepth \maxdimen
\scrollmode

\setbox0\vbox{%
\hsize 6pc
The cat sat on the mat.\par
\showlists
}
```

The annotated output of the `grep` is:

```
# This is an annotation.
# Start of the first line of paragraph.
\hbox(6.94444+0.0)x72.0, glue set 0.58331
# indentation box, 20pt wide
.\hbox(0.0+0.0)x20.0
# The word "The".
.\tenrm T
.\tenrm h
.\tenrm e
# normal interword glue
.\glue 3.33333 plus 1.66666 minus 1.11111
# The word "cat".
.\tenrm c
.\tenrm a
.\tenrm t
.\glue 3.33333 plus 1.66666 minus 1.11111
.\tenrm s
.\tenrm a
.\tenrm t
# Zero width line filling glue.
.\glue(\rightskip) 0.0
# Penalty for breaking page at this point.
\penalty 300
```

```
# Interline glue.
\glue(\baselineskip) 5.05556
# Start of second line of paragraph.
\hbox(6.94444+0.0)x72.0, glue set 20.88878fil
.\tenrm o
.\tenrm n
.\glue 3.33333 plus 1.66666 minus 1.11111
.\tenrm t
.\tenrm h
.\tenrm e
.\glue 3.33333 plus 1.66666 minus 1.11111
.\tenrm m
.\tenrm a
.\tenrm t
.\tenrm .
# Inserted by TeX, for internal reasons.
.\penalty 10000
# Allow the last line of para to be short.
.\glue(\parfillskip) 0.0 plus 1.0fil
.\glue(\rightskip) 0.0
```

This is a complete description of the paragraph formed by T_EX's line breaking algorithm.

This is the essence of the interface between Python and T_EX. Material is sent to T_EX to be typeset, say using `stdin`. The `\showlists` command is used to write the results to `stdout`, from which they are picked up by Python.

On the input side, Python is responsible for parsing the input stream, and placing appropriate items on the horizontal list. It is also responsible for ensuring that nothing inappropriate is placed on the list. The whole process should not generate T_EX errors (although warnings about overfull boxes and so forth are welcome).

On the output side, Python parses the output stream, and from it reconstitutes the boxes that T_EX has formed, thus forming a Python object.

```
In Python code, our example might look like
hlist = Hlist() % new horizontal list
text = "The cat sat on the mat."
hlist.extend(text)
vlist = hlist.linebreak(hsize=6*pica)
```

where hidden in the call to the `linebreak()` method there is a sending of data to T_EX, and a reconstruction in Python of the boxes that T_EX built. From here on, the output or page-makeup routine can be written in Python. Note that `cat_sat.tex` invokes no T_EX macros.

Conclusions

In the 15 years since T_EX has been frozen, very few deficiencies have been found in the algorithms it uses for breaking a paragraph into lines, for typesetting mathematics, and for setting tables. Since 1990 there has been little (but valuable) progress in

the area of Unicode fonts, for which an extension of T_EX genuinely is needed. T_EX was written to be archival, and it is holding up well after its first quarter-century or so.

There are many problems in our use of T_EX. This paper has discussed several:

- coloured text and other decorations,
- interactive use of T_EX,
- input transformation,
- programming page makeup.

All of these arise not out of T_EX itself, but out of the way in which we use T_EX.

There is an irony in our use of T_EX macros. Recall that when Don was looking at the design of T_EX he found that: [11, p.648]:

Every system I looked at had its own universal Turing machine built into it somehow, and everybody’s machine was a little different from everybody else’s.

He then went on to say:

I was tired of having to learn yet another almost-the-same programming language for every system I looked at; I was going to try to avoid that.

What we have now with T_EX macros is a Turing machine very different from any other. This is just what he wished to avoid. However, Q_AT_EX provides a powerful complement to existing T_EX macro packages, and PyT_EX will use T_EX as “just a typesetting language”, which is what Don wanted it to be in the first place.

In 1996, Piet van Oostrum asked Don Knuth about T_EX’s macro programming language [11, p648–9]:

I don’t know if you have ever looked into the L^AT_EX code inside, but if you look into that, you get the impression that T_EX is not the most appropriate programming language to design such a large system. Did you ever think of T_EX being used to program such large systems and if not, would you think of giving it a better programming language?

In response to this, Don Knuth said (loc. cit.):

It would be nice if there were a well-understood standard for an interpretive programming language inside an arbitrary application. Take regular expressions—I define UNIX as “30 definitions of regular expressions living under one roof.” [laughter] Every part of UNIX has a slightly different regular expression. Now, if there were a

universal simple interpretive language that was common to other systems, naturally I would have latched onto that right away.

The theme of this paper is T_EX typesetting with fewer macros. We use instead a “simple interpretive language”, namely Python. If we learn to use T_EX in new ways, and take good care of it, T_EX will be good for its second quarter-century.

Long live T_EX!

References

- [1] eval4tex, <http://www.ccs.neu.edu/home/dorai/tex2page/eval4tex-doc.html>
- [2] Jonathan Fine, Editing .dvi files, or Visual T_EX, *TUGboat*, **17** (1996), 255–259
- [3] ———, Active T_EX and the DOT input syntax, *TUGboat*, **20** (1999), 248–254
- [4] ———, Line breaking and page breaking, *TUGboat*, **21** (2000), 210–221
- [5] ———, Instant Preview and the T_EX daemon, *EuroT_EX 2001 Conference Proceedings*, 49–58
- [6] ———, T_EX as a callable function, *EuroT_EX 2002 Conference Proceedings*, 26–35
- [7] Donald E. Knuth, The Errors of T_EX, *Software — Practice and Experience*, **19** (1989), 607–685. (Reprinted in [9])
- [8] ———, The Future of T_EX and METAFONT, *TUGboat*, **11** (1990), 489 (Reprinted in [13])
- [9] ———, *Literate Programming*, CSLI (1992)
- [10] ———, Questions and Answers II, *TUGboat*, bf 17 (1996), 355–367 (Reprinted in [13])
- [11] ———, Questions and Answers III, *MAPS (Minutes and Appendices)*, **16** 1996, 38–49 (Reprinted in [13])
- [12] ———, The future of T_EX and METAFONT, *TUGboat*, **11** (1990), 489 (reprinted in [13])
- [13] ———, *Digital Typography*, CSLI (1999)
- [14] Donald E. Knuth & Pierre MacKay, Mixing Right-to-Left Texts with Left-to-Right Texts, *TUGboat*, **8**, (1987), 14–25. (Reprinted in [13])
- [15] Frank Mittelbach & Chris Rowley, The L^AT_EX3 Project, <http://www.latex-project.org/guides/ltx3info.pdf>
- [16] NTG T_EX future working group, T_EX in 2003: Part I Propositions and Conjectures on the Future of T_EX, *MAPS (Minutes and Appendices)*, **21** 1998, 13–19
- [17] PerlT_EX, <http://www.ctan.org/tex-archive/macros/latex/contrib/perltext>
- [18] Pymacs, <http://pymacs.progiciels-bpi.ca>
- [19] sT_EXme, <http://stexme.sourceforge.net>