

---

## Creating accessible pdfs with L<sup>A</sup>T<sub>E</sub>X

Ulrike Fischer

### Abstract

This article describes the current state and planned actions to improve accessibility of pdfs created with L<sup>A</sup>T<sub>E</sub>X, as currently undertaken by the L<sup>A</sup>T<sub>E</sub>X Team.

### 1 Accessibility of pdf

The pdf language is at its core a *page description* programming language. It describes very accurately how text and graphical elements look and where they are placed on a page. But it doesn't describe the semantic meaning of the elements and the reading order; by looking at the code there is no way to know if a text is a section heading or some watermark or a footnote or if it belongs to a tabular. You can't know where a sentence has its continuation on another page or how many words a text contains, and sometimes it is even impossible to identify the characters; you only see some glyph index number and copy & paste can give gibberish.

All this is not a problem as long as the pdf is merely meant for printing or viewing but it restricts its use for digital processing like copy & paste, automatic extraction of billing data, reflowing, or using the pdf with a screen reader. For such uses you need accessible, structured, extractable content.

“Accessibility” as a standard is specified in PDF/UA. It is also included in other standards, notably PDF/A-1a and PDF/A-2a (the “a” stands for accessible). The standards contain a number of requirements to help in retrieving the content for further processing; for example, that every character has a Unicode representation (no gibberish when copy & pasting), that word spaces are correctly marked up (so that reflowing works), that the language of the document and the text is declared (so a screen reader can guess the pronunciation), that pictures have sensible alternative descriptions. And most importantly: that the document is *tagged*. This last requirement is responsible for adding structure information to the content. It marks up content as section or tabular cell or list item. This improves navigation in the document with, for example, a screen reader, but also exporting to other formats like XML.

TUG maintains a web page with links to relevant standards, articles and packages [3].

### 2 Creating a tagged pdf

*Tagging* consists of two main tasks. First, in the *stream object* of a page every bit of content must be marked and labelled with a number MCID  $n$ .

The following listing shows a small example. The BDC and the EMC lines are the start and end markers needed for tagging. The /H1 indicates that the content is part of a sectioning element.

```
stream
/H1 <<MCID 0>> BDC
BT
/F17 14.3462 Tf 124.802 706.129
  Td [(0.1)-1100(Section)]TJ
ET
EMC
```

In the second step a number of pdf objects must be created to describe the *structure tree*. Every object contains references to the parent /P and to one or more kid elements /K. The leaf nodes are the MCID  $n$  created in the first step. A typical object looks roughly like this:

```
5 0 obj
<<
  /Type /StructElem
  /S /H1
  /P 4 0 R
  /K <</Type /MCR /MCID 0>>
>>
endobj
```

This is a structure element of the type /H1 (and so a sectioning element) and it has one kid element, the text of the section marked above.

Beside this a number of additional settings and objects must be added to the pdf for cross-referencing and “administration”.

### 3 Changing L<sup>A</sup>T<sub>E</sub>X

Measured in computer time, L<sup>A</sup>T<sub>E</sub>X is quite old. L<sup>A</sup>T<sub>E</sub>X is not only a format; it was always meant to be extended by packages and classes, and over time, many people have contributed to L<sup>A</sup>T<sub>E</sub>X. It has a quite large user base with *very* varied demands regarding stability, features and development. L<sup>A</sup>T<sub>E</sub>X is still used with a variety of engines: pdfT<sub>E</sub>X, X<sub>Y</sub>T<sub>E</sub>X, LuaT<sub>E</sub>X, (u)pT<sub>E</sub>X and backends (dvips, dvi<sub>ps</sub>dmx). One could compare L<sup>A</sup>T<sub>E</sub>X to an old city; lots of houses built at different times in different styles by various people, some modern, some older, some are in a good state, others are falling apart but nevertheless home to someone.

This means that changing L<sup>A</sup>T<sub>E</sub>X is not easy; we can't break lots of packages and old documents even if the reward is accessible pdfs. And we have to consider that documents must be compilable in T<sub>E</sub>X

---

First published in *ArsT<sub>E</sub>Xnica* 28 (Oct 2019), pp. 135–139. Reprinted with permission.

systems of varying age, for example when uploading them to a journal.

Thus, a very important aspect of the project is to develop a long term change strategy and manage integration of core support across the L<sup>A</sup>T<sub>E</sub>X universe.

#### 4 First steps towards tagging

Tagging pdf with L<sup>A</sup>T<sub>E</sub>X has been on the agenda for quite some time. Babet Schalit wrote a thesis about it in 2007, and Ross Moore has given a number of talks and articles at TUG conferences since then. When I considered working on the topic some time ago I got code from both, and decided rather quickly that first some work on the basics was needed.

Tagging should, in my opinion, not be done by creating a package that patches all sorts of commands in other packages; this is much too fragile. It needs proper support in the L<sup>A</sup>T<sub>E</sub>X kernel and proper support in the main classes and packages. I also thought that to identify the needed support and to test implementations and interfaces, concrete code was needed. So I wrote the package `tagpdf`. The package offers core commands to tag a document and to activate some of the other requirements needed to make a pdf accessible. The low-level code to mark up a text as a section looks roughly like this:

```
\tagstructbegin{tag=H1}
  \tagmcbegin{tag=H1}
    Section
  \tagmcbend
\tagstructend
```

The `\tagstructXX` commands create the structure, while the `\tagmcXX` commands add the MCID marks to the page stream.

At present, the `tagpdf` package works with pdfL<sup>A</sup>T<sub>E</sub>X and LuaL<sup>A</sup>T<sub>E</sub>X, and LuaL<sup>A</sup>T<sub>E</sub>X gives the best results, as one doesn't have to worry about the behaviour at page breaks. However, with the help of the work on the `pdfresources` project described below, it should be possible to extend it to other engines and backends.

#### 5 L<sup>A</sup>T<sub>E</sub>X-dev

Another important step towards accessible pdfs was the implementation of the latex-dev format by the L<sup>A</sup>T<sub>E</sub>X team and the maintainers of T<sub>E</sub>X Live and MiK<sub>T</sub><sub>E</sub>X; latex-dev is a pre-release of L<sup>A</sup>T<sub>E</sub>X from the development branch and made available through CTAN. It allows users of a current T<sub>E</sub>X distribution to test their documents and code against an upcoming L<sup>A</sup>T<sub>E</sub>X release by simply using their preferred latex with the suffix `-dev`; `pdflatex-dev`, `lualatex-dev`, `xelatex-dev`, etc. (For more information, see [1].)

latex-dev has not been created solely with tagging in mind, but it will help us to coordinate and test changes with package and class authors, so it is an important part of the project.

#### 6 PDF resource management

When tagging a pdf one has to add a number of settings to pdf dictionaries which can be described as “global resources”. As already mentioned in an answer [2], L<sup>A</sup>T<sub>E</sub>X has no interfaces for this:

Unhappily, the L<sup>A</sup>T<sub>E</sub>X format has overslept PDF development quite entirely. Managing global resources is the prime task for an OS, format in T<sub>E</sub>X speak. Because of the missing resource manager, both [`tikz` and `transparent`] packages do what most packages do, they think they are alone and add their stuff to the resource, ...

With tagging entering the scene it was clear that something needed to be done to remedy this problem and so the `pdfresources` project in the L<sup>A</sup>T<sub>E</sub>X github was created; it contains a (still quite experimental) `expl3`-style interface which offers commands to add contents to pdf resources in a controlled way. It also offers backend-independent interfaces to a number of core commands needed when writing objects to a pdf. The package works with the main engines (pdfT<sub>E</sub>X, LuaT<sub>E</sub>X and XeT<sub>E</sub>X) and backends (dvi<sub>pdf</sub>mx and — more or less — dvips).

The main task for the next months is to test the code, to integrate it into the kernel and to adapt existing packages to use it. The number of packages which should use the pdf resource manager is not very large but includes important packages such as `hyperref`, `tikz`, `media9`, `pdfx`.

#### 7 Adapting the engines

Another open issue that emerged during the last year was missing functionality in engines and backends. For example pdfT<sub>E</sub>X was not ready for pdf 2.0; it has no command to set a major pdf version. (This will be remedied in T<sub>E</sub>X Live 2020.) As pdf 2.0 adds important features needed for accessibility (the concept of associated files) this is clearly something that should be changed. It would be also useful if pdfT<sub>E</sub>X could execute code at shipout time as can be done with luatex with `\latelua`. The dvi<sub>pdf</sub>mx backend and dvips are missing additional color stacks.

#### 8 Adding hooks

As already shown in sections 2 and 4, tagging a pdf requires adding quite a number of commands. Obviously, all the standard structures should if possible add the needed code automatically. For this, hooks

are needed at the right places. The “right place” has firstly a technical meaning; with the exception of LuaTeX, the tagging code inserts *whatsits*; this means it can change the output if used in the wrong place (as sometimes anchors set by `hyperref` do).

But more importantly the “right place” means that we need to identify the *owner* of the code which should insert the tagging code. For example, sections are generally created with `\@startsection`. So this kernel command looks like a natural place to insert hooks for tagging commands. On the other hand, chapters and parts have special commands created by the classes. Does it make sense if the kernel handles the one part and the classes the other? Other examples are bibliographies and glossaries; packages like `biblatex` and `glossaries` look like the natural owner here — and both packages already have lots of hooks which make it easy to implement tagging — but both also use standard structures like lists or tabulars and additions to these generic environments could clash with their needs.

This means that besides a pdf resource manager we also need a hook management. And we need lots of real use cases and examples to be able to investigate the various dependencies.

## 9 Mathematics

How to tag maths is still an open problem. There are quite a number of possibilities to make it accessible.

- One is to attach the L<sup>A</sup>T<sub>E</sub>X source code, either as file or verbatim, with `/ActualText` to the math structure. For a number of environments this can be automated quite well as the `axessibility` package demonstrates (but it is difficult for inline math input with `$. . . $`). The usability with a screen reader is not bad — even if not every word was correctly read aloud in my tests — but it requires that the user understands L<sup>A</sup>T<sub>E</sub>X input syntax, and with large equations and complicated grouping it can be quite difficult to follow and to navigate through subequations. Usability can be improved if one invests the time to manually split the math and add explaining words.
- Another possibility is to mark all the maths bits with MathML structure names. At least with LuaTeX this can probably be done more or less automatically — proofs of concept are the ConTeXt format and TeX4ht. But it is unknown whether screen readers or other applications can actually use the information.
- A third possibility is to convert the equation to MathML, for example with MathJax, and attach it as an associated file to the structure.

But here too it is unclear how such MathML can be processed by the pdf consumer. It is also unknown whether the presentation or content flavour of MathML should be used in this case.

The pdf standard requires that glyphs and symbols are mapped to Unicode. Here too variants are possible; for example, *a* could be mapped to U+1D44E (Mathematical Italic Small A) or U+0061 (Latin Small Letter A);  $\int$  could be mapped to U+222B (Integral) or to `\int` (as is done by the package `mmap`). The first alternative sounds more Unicode-like but actually the screen readers don’t seem to know what to do with the symbols.

The main task here is to get more information to be able to decide about which route to follow.

## 10 Contacts

Quite a number of questions and projects circle around the pdf specification, the needs of users and of pdf consumer applications. To get tagging working it is not enough to know how TeX works. So one important part of the tagging project is to get in contact with people having inside knowledge about pdf and pdf consumer applications in various pdf related organizations and to promote the project in the TeX world to get user feedback.

## 11 Summary

Adding tagging facilities to L<sup>A</sup>T<sub>E</sub>X is a large project with many aspects. Happily it doesn’t have to be done in one large jump; with the `tagpdf` package it is already possible for adventurous users with a bit of knowledge in TeX programming to tag quite large documents. Despite the clear warning in the documentation that it isn’t meant for production, I have already received feedback about several successful uses. This gives hope that it can evolve to a stable and usable system.

## References

- [1] L<sup>A</sup>T<sub>E</sub>X. L<sup>A</sup>T<sub>E</sub>X development formats are now available, 2019. <https://latex-project.org/news/2019/09/01/LaTeX-dev-format/>
- [2] H. Oberdiek. `TikZ` and `transparent` incompatibility, 2015. <https://tex.stackexchange.com/a/253417/2388>
- [3] TUG. PDF accessibility and PDF standards, 2019. <https://tug.org/twg/accessibility/>

◇ Ulrike Fischer  
L<sup>A</sup>T<sub>E</sub>X Project  
Mönchengladbach, Germany  
[ulrike.fischer@latex-project.org](mailto:ulrike.fischer@latex-project.org)