

The tagpdf package, v0.80

Ulrike Fischer*

2021-02-23

This package is not meant for normal document production.
You need a very current \LaTeX format. You need a very current L³ programming layer.
You need the new \LaTeX PDF management bundle.
This package is incomplete, experimental and quite probably contains bugs. At some time it will disappear when the code has been integrated into the \LaTeX format.
This package can change in an incompatible way.
You need some knowledge about \TeX , PDF and perhaps even lua to use it.
Issues, comments, suggestions should be added as issues to the github tracker:

<https://github.com/u-fischer/tagpdf>

Contents

1. Preface to version 0.8	2
2. Introduction	3
2.1. Tagging and accessibility	4
2.2. Engines and modes	4
2.3. References	4
2.4. Validation	4
2.5. Examples wanted!	5
2.6. Changes in 0.3	5
2.7. Changes in 0.5	5
2.8. Changes in 0.6	6
2.9. Proof of concept: the tagging of the documentation itself	6
2.10. Changes in version 0.61	7
2.11. Changes in version 0.8	7
3. Setup	8
3.1. Modes and package options	8
3.2. Setup and activation	8

*fischer@troubleshooting-tex.de

4. Tagging	9
4.1. Three tasks	10
4.2. Task 1: Marking the chunks: the mark-content-step	10
4.2.1. Generic mode versus lua mode in the mc-task	13
4.2.2. Commands to mark content and chunks	13
4.2.3. Tips	16
4.2.4. Links	16
4.2.5. Math	18
4.2.6. Split paragraphs	19
4.3. Task 2: Marking the structure	19
4.3.1. Structure types	19
4.3.2. Sectioning	20
4.3.3. Commands to define the structure	20
4.3.4. Root structure	22
4.3.5. Attributes and attribute classes	23
4.4. Task 3: tree Management	23
4.5. A fully marked up document body	24
4.6. Lazy and automatic tagging	25
4.7. Adding tagging to commands	25
5. Alternative text, ActualText and text-to-speech software	26
6. Standard type and new tags	26
7. “Real” space glyphs	27
8. Accessibility is not only tagging	27
9. To-do	28
References	29
A. Some remarks about the PDF syntax	29

1. Preface to version 0.8

⁸³With this version one major step towards integration of the code into the \LaTeX has been done: The code now relies on the new \LaTeX PDF management. This management, which is for a testphase provided as an external package, `pdfmanagement-testphase`, prepares the ground for better support for tagged PDF in \LaTeX . It is part of a larger project to automatically generate tagged PDF <https://www.latex-project.org/news/2020/11/30/tagged-pdf-FS-study/>⁸³

⁸⁴While this is a major improvement—it will for example allow to use `tagpdf` with more engines as the new PDF management supports all major engines and allowed to add support for associated files—it also means that this version requires a special setup of the document and is incompatible with a number of packages, see the documentation of `pdfmanagement-testphase` for details.⁸⁴

⁸⁵Another important step are the new hook management in \LaTeX : the newest development version has hooks for paragraphs which should at the end allow to tag many paragraphs automatically. The small red numbers in this section show them in action. The main problem here is not to tag a paragraph, but to avoid to tag too many: paragraphs pop up in many places. ⁸⁵

2. Introduction

Since many year the creation of accessible PDF-files with \LaTeX which conform to the PDF/UA standard has been on the agenda of \TeX -meetings. Many people agree that this is important and Ross Moore has done quite some work on it. There is also a TUG-mailing list and a webpage [2] dedicated to this theme.

But in my opinion missing are means to *experiment* with tagging and accessibility. Means to try out, how difficult it is to tag some structures, means to try out, how much tagging is really needed (standards and validators don't need to be right ...), means to test what else is needed so that a PDF works e.g. with a screen reader. Without such experiments it is imho quite difficult to get a feeling about what has to be done, which kernel changes are needed, how packages should be adapted.

This package tries to close this gap by offering *core* commands to tag a PDF¹

My hope is that the knowledge gained by the use of this package will at the end allow to decide if and how code to do tagging should be part of the \LaTeX kernel.

The package does not patch commands from other packages. It is also not an aim of the package to develop such patches. While at the end changes to various commands in many classes and packages will be needed to get tagged PDF files – and the examples accompanying the package try (or will try) to show various strategies – these changes should in my opinion be done by the class, package and document writers themselves using a sensible API provided by the kernel and not by some external package that adds patches everywhere and would need constant maintenance – one only need to look at packages like tex4ht or bidi or hyperref to see how difficult and sometimes fragile this is.

So this package deliberately concentrates on the basics – and this already quite a lot, there are much more details involved as I expected when I started.

I'm sure that it has bugs. Bugs reports, suggestions and comments can be added to the issue tracker on github. <https://github.com/u-fischer/tagpdf>.

Please also check the github site for new examples and improvements.

¹In case you don't know what this means: there will be some explanations later on.

2.1. Tagging and accessibility

While the package is named `tagpdf` the goal is actually *accessible* PDF-files. Tagging is *one* requirement for accessibility but there are others. I will mention some later on in this documentation, and – if sensible – I will also try to add code, keys or tips for them.

So the name of the package is a bit wrong. As excuse I can only say that it is shorter and easier to pronounce.

2.2. Engines and modes

The package works currently with `pdflatex` and `lualatex`. First steps have been done to also enable support for `xelatex` and the `latex-dvips-route`; but this isn't yet much tested.

The package has two modes: the *generic mode* which should work in theory with every engine and the *lua mode* which works only with `lualatex`.

I implemented the generic mode first. Mostly because my tex skills are much better than my lua skills and I wanted to get the tex side right before starting to fight with attributes and node traversing.

While the generic mode is not bad and I spent quite some time to get it working I nevertheless think that the lua mode is the future and the only one that will be usable for larger documents. PDF is a page orientated format and so the ability of `luatex` to manipulate pages and nodes after the \TeX -processing is really useful here. Also with `luatex` characters are normally already given as unicode. The main problem with `luatex` is how to insert “fake spaces” between words.

2.3. References

My main reference was the free reference for PDF 1.7. [1]. This document is from 2006.

In the meantime PDF 2.0. has been released. I know that it contains also for accessibility relevant changes. As I got now the specification and `luatex` can set with `\pdfmajorversion` the version to 2.0 I will start to check to look into it too.

2.4. Validation

PDF's created with the commands of this package must be validated:

- One must check that the PDF is *syntactically* correct. It is rather easy to create broken PDF: e.g. if a chunk is opened on one page but closed on the next page.
- One must check how good the requirements of the PDF/UA standard are followed *formally*.
- One must check how good the accessibility is *practically*.

Syntax validation and formal standard validation can be done with preflight of the (non-free) adobe acrobat. It can also be done also with the free PDF Accessibility Checker (PAC 3) [4].

Syntax validation and formal standard validation can be done with preflight of the (non-free) adobe acrobat. It can also be done also with the free PDF Accessibility Checker (PAC 3) [4]. There is also the validator veraPDF [3]. But I didn't try it yet and have no idea if it is useful here.

Practical validation is naturally the more complicated part. It needs screen reader, users which actually knows how to handle them, can test documents and can report where a PDF has real accessibility problems.

Preflight woes

Sadly validators can not be always trusted. As an example for an reason that I don't understand the adobe preflight don't like the list structure L. It is also possible that validators contradict: that the one says everything is okay, while the other complains.

2.5. Examples wanted!

To make the package usable examples are needed: example that demonstrates how various structures can be tagged and which patches are needed, examples for the test suite, examples that demonstrates problems.

Feedback, contributions and corrections are welcome!

All examples should use the tagpdfsetup key `uncompress` described in the next section so that uncompressed PDF are created and the internal objects and structures can be inspected and – hopefully soon – be compared by the l3build checks.

2.6. Changes in 0.3

In this version I improved the handling of alternative and actual text. See section 5. This change means that the package now relies on the experimental expl3 package `l3str-convert`.

I no longer try to (pdf-)escape the tag names: it is a bit unclear how to do it at best with luatex. This will perhaps later change again.

2.7. Changes in 0.5

I added code to handle attributes and attribute classes, see section 4.3.5 and corrected a small number of code errors.

I added code to add “real” space glyphs to the PDF, see section 7.

2.8. Changes in 0.6













Breaking change! The attributes used in luamode to mark the MC-chunks are no longer set globally. I thought that global attribute would make it easier to tag, but it only leads to problem when e.g. header and footer are inserted. So from this version on the attributes are set locally and the effect of a `\tagmcbegin` ends with the current group. This means that in some cases more `\tagmcbegin` are needed and this affected some of the examples, e.g. the patching commands for sections with KOMA. On the other side it means that quite often one can omit the `\tagmcbegin` command.

2.9. Proof of concept: the tagging of the documentation itself

With version 0.6 the documentation itself has been tagged. The tagging is in no way perfect. The validator from Adobe doesn't complain, but PAX3 wanted alternative text for all links (no idea why) and so I put everywhere simple text like "link" and "ref". The links to footnote gave warnings, so I disabled them. I used types from the PDF version 1.7, mostly as I have no idea what should be used for code in 2.0. Margin notes are probably simply wrong ...

But even if the documentation passed the tests of the validators: as mentioned above passing a formal test doesn't mean that the content is really good and usable. I have a lot doubts that the code parts are really readable. The bibliography and the references must be improved. The user commands used for the tagging and also some of the patches used are rather crude. So there is lot space for improvement.

Be aware that to create the tagged version a very experimental package and a experimental hyperref driver was needed for the links. This files are not in the package!

Dateiname tagpdf.pdf			
Sprache en-UK		Tags 814	Seiten 29
PDF/UA			
 Die von PAC geprüften PDF/UA-Anforderungen sind erfüllt			
Prüfpunkt	Bestan...	Warnung	
 PDF Syntax	65	0	
 Schriften	18	0	
 Inhalt	72892	0	
 Eingebettete Dateien	0	0	
 Natürliche Sprache	35008	0	
 Strukturelemente	296	0	
 Strukturbaum	16	0	
 Rollenzuordnungen	22	0	
 Alternative Beschreibungen	32	0	
 Metadaten	0	0	
 Dokumenteinstellungen	34	0	

2.10. Changes in version 0.61

- internal code adaption to expl3 changes.
- dropped the compresslevel key – probably not needed

2.11. Changes in version 0.8

As a first step to include the code proper in the \LaTeX kernel the module name has changed from `uftag` to `tag`. The commands starting with `\uftag` will stay valid for some time but then be deprecated.

Breaking change! The argument of `newattribute` option should no longer add the dictionary bracket `<<...>>`, they are added by the code.

Breaking change! The package now requires the new PDF management as provided for now by the package `pdfmanagement-testphase`

Support to add associated files to structures has been added with new keys `AF`, `AFinline` and `AFinline-o`.

Breaking change! The support for other 8-bit input encodings has been removed. `utf8` is now the required encoding.

The keys `|lang|`, `|ref|` and `|E|` have been added for structures.

3. Setup

Activation needed!

When the package is loaded it will – apart from loading more packages and defining a lot of things – not do anything. You will have to activate it with `\tagpdfsetup`, see below. (At least that's the theory, I'm not quite sure, if really the tests are done always as planned....)

3.1. Modes and package options

The package has two different modes: The **generic mode** works (in theory, currently only with pdfTeX and LuaTeX) probably with all engines, the **lua mode** only with LuaTeX. The differences between both modes will be described later. The mode can be set with package options:

`luamode`

This is the default mode. It will use the generic mode if the document is processed with pdfLaTeX and the lua mode with LuaLaTeX.

`genericmode`

This will force the generic mode for all engines.

3.2. Setup and activation

The following command setups the general behaviour of the package. The command should be normally used only in the preamble (for a few keys it could also make sense to change them in the document).

`\tagpdfsetup{<key-val-list>}`

The key-val list understands the following keys:

activate-all Boolean, initially false. Activates everything, that's normally the sensible thing to do.

activate-mc Boolean, initially false. Activates the code related to marked content.

activate-struct Boolean, initially false. Activates the code related to structures. Should be used only if **activate-mc** has been used too.

activate-tree Boolean, initially false. Activates the code related to trees. Should be used only if the two other keys have been used too.

add-new-tag Allows to define new tag names, see section 6 for a description.

check-tags Boolean, initially true. Activates some safety checks (but doesn't do very much currently. It will perhaps be merged with the log-level key).

interwordspace Choice key, possible values are **true**/**on** and **false**/**off**. The key activates/deactivates the insertion of space glyphs, see section 7. In the luamode it only works if at least **activate-mc** has been used.

log Choice key, possible values **none**, **v**, **vv**, **vvv**, **all**. Setups the log level. Changing the value affects currently mostly the luamode: "higher" values gives more messages in the log. The current levels and messages have been setup in a quite ad-hoc manner and will need improvement.

newattribute This key takes two arguments and declares an attribute. See 4.3.5.

luamode **show-spaces** Boolean. That's a debug option, it helps to see where space glyph will be inserted if **interwordspace** is activated.

tabsorder Choice key, possible values are **row**, **column**, **structure**, **none**. This decides if a /Tabs value is written to the dictionary of the page objects. Not really needed for tagging itself, but one of the things you probably need for accessibility checks. So I added it. Currently the tabsorder is the same for all pages. Perhaps this should be changed

luamode **tagunmarked** Boolean, initially true. When this boolean is true, the lua code will try to mark everything that has not been marked yet as an artifact. The benefit is that one doesn't have to mark up every deco rule oneself. The danger is that it perhaps marks things that shouldn't be marked – it hasn't been tested yet with complicated documents containing annotations etc. See also section 4.6 for a discussion about automatic tagging.

uncompress Sets both the PDF compresslevel and the PDF objcompresslevel to 0 and so allows to inspect the PDF.

4. Tagging

pdf is a page orientated graphic format. It simply puts ink and glyphs at various coordinates on a page. A simple stream of a page can look like this²:

```
stream
  BT
  /F27 14.3462 Tf           %select font
  89.291 746.742 Td        %move point
  [(1)-574(Intro)-32(duction)]TJ %print text
  /F24 10.9091 Tf         %select font
  0 -24.35 Td              %move point
  [(Let's)-331(start)]TJ   %print text
  205.635 -605.688 Td      %move point
  [(1)]TJ                  %print text
```

²The appendix contains some remarks about the syntax of a PDF file

```
ET
endstream
```

From this stream one can extract the characters and their placement on the page but not their semantic meaning (the first line is actually a section heading, the last the page number). And while in the example the order is correct there is actually no guaranty that the stream contains the text in the order it should be read.

Tagging means to enrich the PDF with information about the *semantic* meaning and the *reading order*. (Tagging can do more, one can also store all sorts of layout information like font properties and indentation with tags. But as I already wrote this package concentrates on the part of tagging that is needed to improve accessibility.)

4.1. Three tasks

To tag a PDF three tasks must be carried out:

- | | |
|-------------|---|
| mc-task | 1. The mark-content-task: The document must add “labels” to the page stream which allows to identify and reference the various chunks of text and other content. This is the most difficult part of tagging – both for the document writer but also for the package code. At first there can be quite many chunks as every one is a leaf node of the structure and so often a rather small unit. At second the chunks must be defined page-wise – and this is not easy when you don’t know where the page breaks are. Also in a standard document a lot text is created automatically, e.g. the toc, references, citations, list numbers etc and it is not always easy to mark them correctly. |
| struct-task | 2. The structure-task: The document must declare the structure. This means marking the start and end of semantically connected portions of the document (correctly nested as a tree). This too means some work for the document writer, but less than for the mc-task: at first quite often the mc-task and the structure-task can be combined, e.g. when you mark up a list number or a tabular cell or a section header; at second one doesn’t have to worry about page breaks so quite often one can patch standard environments to declare the structure. On the other side a number of structures end in \LaTeX only implicitly – e.g. an item ends at the next item, so getting the PDF structure right still means that additional mark up must be added. |
| tree-task | 3. The tree management: At last the structure must be written into the PDF. For every structure an object of type <code>StructElem</code> must be created and flushed with keys for the parents and the kids. A <code>parenttree</code> must be created to get a reference from the mc-chunks to the parent structure. A <code>rolemap</code> must be written. And a number of dictionary entries. All this is hopefully done automatically and correctly by the package |

4.2. Task 1: Marking the chunks: the mark-content-step

To be able to refer to parts of the text in the structure, the text in the page stream must get “labels”. In the PDF reference they are called “marked content”. The three main variants needed here are:

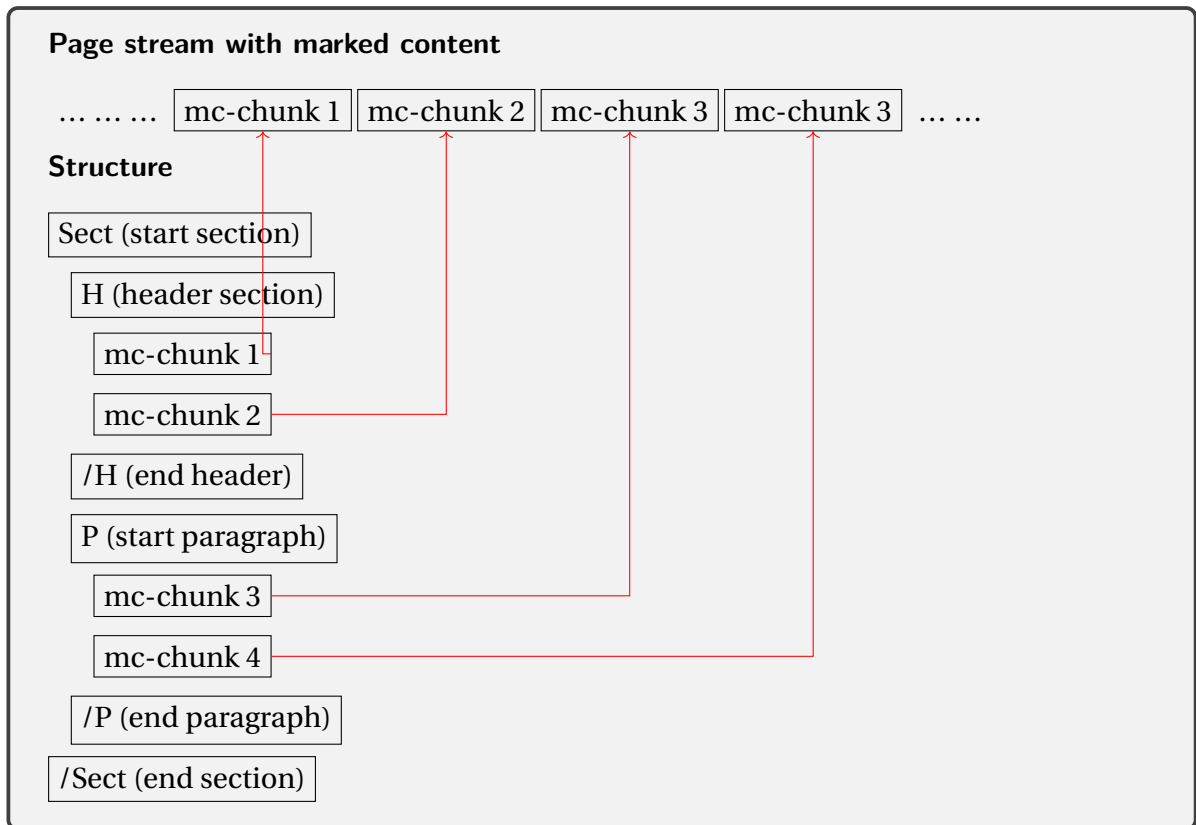


Figure 1: Schematical description of the relation between marked content in the page stream and the structure

Artifacts They are marked with of a pair of keywords, BMC and EMC which surrounds the text. BMC has a single prefix argument, the fix tag name `/Artifact`. Artifacts should be used for irrelevant text and page content that should be ignored in the structure. Sadly it is often not possible to leave such text simply unmarked – the accessibility tests in Acrobat and other validators complain.

```
/Artifact BMC
text to be marked
/EMC
```

Artifacts with a type They are marked with of a pair of keywords, BDC and EMC which surrounds the text. BDC has two arguments: again the tag name `/Artifact` and a following dictionary which allows to specify the suppressed info. Text in header and footer can e.g. be declared as pagination like this:

```
/Artifact <</Type /Pagination>> BDC
text to be marked
/EMC
```

Content Content is marked also with of a pair of keywords, BDC and EMC. The first argument

of BDC is a tag name which describes the structural type of the text³ Examples are /P (paragraph), /H2 (header), /TD (table cell). The reference mentions a number of standard types but it is possible to add more or to use different names.

In the second argument of BDC – in the property dictionary – more data can be stored. *Required* is an /MCID-key which takes an integer as a value:

```
/H <</MCID 3>> BDC
  text to be marked
/EMC
```

This integer is used to identify the chunk when building the structure tree. The chunks are numbered by page starting with 0. As the numbers are also used as an index in an array they shouldn't be "holes" in the numbering system (It is perhaps possible to handle a numbering scheme not starting by 0 and having holes, but it will enlarge the PDF as one would need dummy objects.).

It is possible to add more entries to the property dictionary, e.g. a title, alternative text or a local language setting.

The needed markers can be added with low level code e.g. like this (in pdftex syntax):

```
\pdfliteral page {/H <</MCID 3>> BDC}%
  text to be marked
\pdfliteral page {EMC}%
```

This sounds easy. But there are quite a number of traps.

1. PDF is a page oriented format. And this means that the start BDC/BMC and the corresponding end EMC must be on the same page. So marking e.g. a section title like in the following example won't always work as the literal before the section could end on the previous page:

```
\pdfliteral page {/H <</MCID 3>> BDC} %problem: possible pagebreak here
  \section{mysection}
\pdfliteral page {EMC}%
```

Using the literals *inside* the section argument is better, but then one has to take care that they don't wander into the header and the toc.

2. Literals are "whatsits" nodes and can change spacing, page and line breaking. The literal *behind* the section in the previous example could e.g. lead to a lonely section title at the end of the page.
3. The /MCID numbers must be unique on a page. So you can't use the literal in a saved box that you reuse in various places. This is e. g. a problem with `longtable` as it saves the table header and footer in a box.
4. The /MCID-chunks are leaf nodes in the structure tree, so they shouldn't be nested.

³There is quite some redundancy in the specification here. The structural type is also set in the structure tree. One wonders if it isn't enough to use always /SPAN here.

5. Often text in a document is created automatically or moved around: entries in the table of contents, index, bibliography and more. To mark these text chunks correctly one has to analyze the code creating such content to find suitable places to inject the literals.
6. The literals are inserted directly and not at shipout. This means that due to the asynchronous page breaking of T_EX the MCID-number can be wrong even if the counter is reset at every page (this package uses in generic mode a label-ref-system to get around this problem. This sadly means that three compilations are needed until everything has settled down).
7. There exist environments which process their content more than once – examples are `align` and `tabularx`. So one has to check for doublettes and holes in the counting system.
8. PDF is a page oriented format. And this means that the start and the end marker must be on the same page ... *so what to do with normal paragraphs that split over pages??*. This question will be discussed in subsection [4.2.6](#).

4.2.1. Generic mode versus lua mode in the mc-task

While in generic mode the commands insert the literals directly and so have all the problems described above the lua mode works quite differently: The tagging commands don't insert literals but set some *attributes* which are attached to all the following nodes. When the page is shipped out some lua code is called which wanders through the shipout box and injects the literals at the places where the attributes changes.

This means that quite a number of problems mentioned above are not relevant for the lua mode:

1. Pagebreaks between start and end of the marker are *not* a problem. So you can mark a complete paragraph. If a pagebreak occur directly after an start marker or before an end marker this can lead to empty chunks in the PDF and so bloat up PDF a bit, but this is imho not really a problem (compared to the size increase by the rest of the tagging).
2. The commands don't insert literals directly and so affect line and page breaking much less.
3. The numbering of the MCID are done at shipout, so no label/ref system is needed.
4. The code can do some marking automatically. Currently everything that has not been marked up by the document is marked as artifact.

4.2.2. Commands to mark content and chunks

Generic mode only	It is vital that the end command is executed on the same page as the begin command. So think carefully how to place them. For strategies how to handle paragraphs that split over pages see subsection 4.2.6 .
-------------------	--

```
\tagmcbegin{<key-val-list>}
\tag_mc_begin:n{<key-val-list>}
```

These commands insert the begin of the marked content code in the PDF. They don't start a paragraph. *They don't start a group*. The user command additionally issues an `\ignorespaces` to suppress spaces after itself. Such markers should not be nested. The command will warn you if this happens.

The key-val list understands the following keys:

tag This is required, unless you use the **artifact** key. The value of the key is normally one of the standard type listed in section 6 (without a slash at the begin, this is added by the code). It is possible to setup new tags, see the same section. The value of the key is expanded, so it can be a command. The expansion is passed unchanged to the PDF, so it should with a starting slash give a valid PDF name (some ascii with numbers like H4 is fine).

artifact This will setup the marked content as an artifact. The key should be used for content that should be ignored. The key can take one of the values **pagination**, **layout**, **page**, **background** and **notype** (this is the default). Text in the header and footer should be marked with **artifact=pagination**.

It is not quite clear if rules and other decorative graphical objects needs to be marked up as artifacts. Acrobat seems not to mind if not, but PAC 3 complained.

The validators complain if some text is not marked up, but it is not quite clear if this is a serious problem.

lua mode
only

The lua mode will mark up everything unmarked as **artifact=notype**. You can suppress this behaviour by setting the tagpdfsetup key **tagunmarked** to false. See section 3.2.

stash Normally marked content will be stored in the “current” structure. This may not be what you want. As an example you may perhaps want to put a marginnote behind or before the paragraph it is in the tex-code. With this boolean key the content is marked but not stored in the kid-key of the current structure.

label This key sets a label by which you can call the marked content later in another structure (if it has been stashed with the previous key). Internally the label name will start with **tagpdf-**.

alttext This key inserts an `/Alt` value in the property dictionary of the BDC operator. See section 5. The value is handled as verbatim string, commands are not expanded.

alttext-o This key inserts an `/Alt` value in the property dictionary of the BDC operator. See section 5. The value is handled as verbatim string like the key **alttext** but expanded once (the `o` refers to the `o` type in expl3). That means that you can do something like this:

```
\newcommand\myalttext{\frac{a}{b}}
\tagmcbegin{tag=P,alttext-o=\myalttext}
```

and it will insert $\frac{a}{b}$ (hex encoded) in the PDF.

actualtext This key inserts an /ActualText value in the property dictionary of the BDC operator. See section 5. The value is handled as verbatim string, commands are not expanded.

actualtext-o This key inserts an /ActualText value in the property dictionary of the BDC operator. See section 5. The value is handled as verbatim string like the key **actualtext** but expanded once (the o refers to the o type in expl3). That means that you can do something like this:

```
\newcommand\myactualtext{X}
\tagmcbegin{tag=P,alttext-o=\myactualtext}
```

and it will insert X (hex encoded) in the PDF.

raw This key allows you to add more entries to the properties dictionary. The value must be correct, low-level PDF. E.g. raw=/Alt (Hello) will insert an alternative Text.

```
\tagmcend
\tag_mc_end:
```

These commands insert the end code of the marked content. They don't end a group and in generic mode it doesn't matter if they are in another group as the starting commands. The user command also issues at first an \unskip. In generic mode both commands check if there has been a begin marker and issue a warning if not. In luamode it is often possible to omit the command, as the effect of the begin command ends with the current group or with a new \tagmcbegin anyway.

luamode

About nesting mc-commands

Attention! As mentioned above in luamode it is often possible to omit the \tagmcend command. This basically means that mc-chunks can be nested. But be aware that this can have surprising side effects on the order in the structure. Consider a code like this:

```
\tagmcbegin{tag=P} Block 1 {\tagmcbegin{tag=P Block 2}} Block 3 \tagmcend
```

Block 3 will here get the same id number as Block 1 and be stored in the structure *before* Block 2. The PDF will record the chunks in the order Block 1 -- Block 3 -- Block 2.

```
\tagmcuse{\labelname}
\tag_mc_use:n{\labelname}
```

These commands allow you to record a marked content that you stashed away into the current structure. Be aware that a marked content can be used only once – the command will warn you if you try to use it a second time.

```
\tagmcfifinTF{<truecode>}{<false>}  
\tag_mc_if_in:TF{<truecode>}{<false>}
```

These commands check if a marked content is currently open and allows you to e.g. add the end marker if yes.

In *generic mode*, where marked content command shouldn't be nested, it works with a global boolean.

In *lua mode* it tests if the mc-attribute is currently unset. You can't test the nesting level with it!

```
\tagpdfget{<keyword>}  
\tag_get:n{<keyword>}
```

These commands give back some variables. Currently the only working key word is mc_tag..

4.2.3. Tips

- Mark commands inside floats should work fine (but need perhaps some compilation rounds in generic mode).
- In case you want to use it inside a `\savebox` (or some command that saves the text internally in a box): If the box is used directly, there is probably no problem. If the use is later, stash the marked content and add the needed `\tagmcuse` directly before oder after the box when you use it.
- Don't use a saved box with markers twice.
- If boxes are unboxed you will have to analyze the PDF to check if everything is ok.
- If you use complicated structures and commands (breakable boxes like the one from `tcolorbox`, `multicol`, many footnotes) you will have to check the PDF.

4.2.4. Links

Link annotations (like all PDF annotations) are associated with a geometric region of the page rather than with a particular object in its content stream. Any connection between the link and the content is based solely on visual appearance rather than on an explicitly specified association.

To connect such a link with the underlying text a specific structure has to be added, see [2](#). When using the low-level commands to create a link, this structure can be created by adding an entry to the `attr` argument of the `startlink` command, and by adding a command of the `tagpdf` package after the `endlink` command, as show in the following listing:


```

18 0 obj    %Link Object
<< /Type /Annot /Subtype/Link
  /Rect [ 196.109 494.573 399.167 506.831 ]
  /StructParent 16
  /A<</Type/Action /S/URI /URI(https://github.com/u-fischer/tagpdf)>>
>>
endobj

19 0 obj    %Object reference (OBJR)
<</Type /OBJR /Obj 18 0 R>>
endobj

17 0 obj    %Structure Element
<< /Type /StructElem /S /Link /P 11 0 R
  /K [<</Type /MCR /Pg 8 0 R /MCID 6>> 19 0 R ]>>
endobj};

5 0 obj    %Parenttree
<< /Nums
  [... 16 17 0 R ...] >>
endobj

```

Figure 2: Structure needed for a link annotation

```

\pdfextension startlink
  attr
  {
    /StructParent \int_use:N\c@g__tag_parenttree_obj_int %<-----
  }
  user {
    /Subtype/Link
    /A
    <<
      /Type/Action
      /S/URI
      /URI(http://www.dante.de)
    >>
  }
  This is a link.
\pdfextension endlink
\__tag_struct_finish_link: %<-----

```

The question is here is how to add this to the hyperref commands without doing a lot patching? This had lead to a new project which tries to move some definitions for links to the \TeX -kernel so that package likes hyperref and tagpdf can coexist and both add their code to links at the same time.

4.2.5. Math

Math is a problem. I have seen an example where *every single symbol* has been marked up with tags from MathML along with an `/ActualText` entry and an entry with alternate text which describes how to read the symbol. The PDF then looked like this

```

/mn <</MCID 6 /ActualText<FEFF0034>/Alt( : open bracket: four )>>BDC
...
/mn <</MCID 7 /ActualText<FEFF0033>/Alt( third s )>>BDC
...
/mo <</MCID 8 /ActualText<FEFF2062>/Alt( times )>>BDC

```

If this is really the way to go one would need some script to add the mark-up as doing it manually is too much work and would make the source unreadable – at least with pdf \LaTeX and the generic mode. In lua mode is it probably possible to hook into the `mlist_to_hlist` callback and add marker automatically.

But I'm not sure that this is the best way to do math. It looks rather odd that a document should have to tell a screen reader in such detail how to read an equation. It would be much more efficient, sensible and flexible if a complete representation of the equation in mathML could be stored in the PDF and the task how to read this aloud delegated to the screen reader. More investigations are needed here.

See also section 5 for some more remarks and tests.

4.2.6. Split paragraphs

Generic mode only A problem are paragraphs with page breaks. As already mentioned the end marker EMC must be added on the same page as the begin marker. But it is in pdf_latex *very* difficult to inject something at the page break automatically. One can manipulate the shipout box to some extend in the output routine, but this is not easy and it gets even more difficult if inserts like footnotes and floats are involved: the end of the paragraph is then somewhere in the middle of the box.

So with pdf_latex in generic mode one currently has to do the splitting manually.

The example mc-manual-para-split demonstrates how this can be done. The general idea is to use \vadjust in the right place:

```
\tagmcbegin{tag=P}
...
fringilla, ligula wisi commodo felis, ut adipiscing felis dui in
enim. Suspendisse malesuada ultrices ante.% page break
\vadjust{\tagmcbegin{tag=P}}
Pellentesque scelerisque
...
sit amet, lacus.\tagmcbegin{tag=P}
```

4.3. Task 2: Marking the structure

The structure is represented in the PDF with a number of objects of type StructElem which build a tree: each of this objects points back to its parent and normally has a number of kid elements, which are either again structure elements or – as leafs of the tree – the marked contents chunks marked up with the tagmc-commands. The root of the tree is the StructTreeRoot.

4.3.1. Structure types

The tree should reflect the *semantic* meaning of the text. That means that the text should be marked as section, list, table head, table cell and so on. A number of standard structure types is predefined, see section 6 but it is allowed to create more. If you want to use types of your own you must declare them. E.g. this declares two new types TAB and FIG and base them on P:

```
\tagpdfsetup{
  add-new-tag = TAB/P,
  add-new-tag = FIG/P,
}
```

4.3.2. Sectioning

The sectioning units can be structured in two ways: a flat, html-like and a more xml-like version. The flat version creates a structure like this:

```
<H1>section header</H1>
<P> text</P>
<H2>subsection header</H2>
...
```

So here the headers are marked according their level with H1, H2, etc.

In the xml-like tree the complete text of a sectioning unit is surrounded with the Sect tag, and all headers with the tag H. Here the nesting defines the level of a sectioning header.

```
<Sect>
  <H>section header</H>
  <P> text</p>
  <Sect>
    <H>subsection header</H>
    ...
  </Sect>
</Sect>
```

The flat version is more \LaTeX -like and it is rather straightforward to patch `\chapter`, `\section` and so on to insert the appropriate H... start and end markers. The xml-like tree is more difficult to automate. If such a tree is wanted I would recommend to use – like the context format – explicit commands to start and end a sectioning unit.

4.3.3. Commands to define the structure

The following commands can be used to define the tree structure:

```
\tagstructbegin{key-val-list}
\tag_struct_begin:n{key-val-list}
```

These commands start a new structure. They don't start a group. They set all their values globally.

The key-val list understands the following keys:

tag This is required. The value of the key is normally one of the standard type listed in section 6. It is possible to setup new tags/types, see the same section.

stash Normally a new structure inserts itself as a kid into the currently active structure. This key prohibits this. The structure is nevertheless from now on “the current active structure” and parent for following marked content and structures.

label This key sets a label by which you can use the structure later in another structure. Internally the label name will start with `tagpdfstruct-`.

alttext This key inserts an `/Alt` value in the dictionary of structure object, see section 5. The value is handled as verbatim string and hex encoded.

alttext-o This key inserts an `/Alt` value in the dictionary of a structure object, see section 5. The value is handled as verbatim string like the key **alttext** but expanded once (the `o` refers to the `o` type in expl3). That means that you can do something like this:

```
\newcommand\myalttext{\frac{a}{b}}
\tagstructbegin{tag=P,alttext-o=\myalttext}
```

and it will insert `\frac{a}{b}` (hex encoded) in the PDF.

actualtext This key inserts an `/ActualText` value in the dictionary of structure object, see section 5. The value is handled as verbatim string, commands are not expanded.

actualtext-o This key inserts an `/ActualText` value in the dictionary of structure object, see section 5. The value is handled as verbatim string like the key **actualtext** but expanded once (the `o` refers to the `o` type in expl3). That means that you can do something like this:

```
\newcommand\myactualtext{X}
\tagstructbegin{tag=P,alttext-o=\myactualtext}
```

and it will insert `X` (hex encoded) in the PDF.

attribute This key takes as argument a comma list of attribute names (use braces to protect the commas from the external key-val parser) and allows to add one or more attribute dictionary entries in the structure object. As an example

```
\tagstructbegin{tag=TH,attribute= TH-row}
```

See also section 4.3.5.

attribute-class This key takes as argument a comma list of attribute names (use braces to protect the commas from the external key-val parser) and allows to add one or more attribute classes to the structure object. As an example

```
\tagstructbegin{tag=TH,attribute-class= TH-row}
```

See also section 4.3.5.

title This key allows to set the dictionary entry `/Title` in the structure object. The value is handled as verbatim string and hex encoded. Commands are not expanded.

title-o This key allows to set the dictionary entry `/Title` in the structure object. The value is expanded once and then handled as verbatim string like the **title** key.

AF This key allows to reference an associated file in the structure element. The value should be the name of an object pointing to the `/Filespec` dictionary as expected by `\pdf_object_ref:n` from a current l3kernel. For example:

```
\pdfdict_put:nnn {l_pdffile/FileSpec} {AFRelationship}{/Supplement}
\pdffile_embed_file:nnn{example-input-file.tex}{\tag/AFtest}
\tagstructbegin{tag=P,AF=tag/AFtest}
```

As shown, the wanted AFRelationship can be set by filling the dictionary with the value. The mime type is here detected automatically, but for unknown types it can be set too. See the `l3pdf` documentation for details. Associated files are a concept new in PDF 2.0, but the code currently doesn't check the pdf version, it is your responsibility to set it (this can be done with the `pdfversion` key in `\DeclareDocumentMetadata`).

AFinline This key allows to embed an associated file with inline content. The value is some text, which is embedded in the PDF as a text file with mime type `text/plain`.

```
\tagstructbegin{tag=P,AFinline=Some extra text}
```

AFinline-o This is like `verb+AFinline+`, but it expands the value once.

lang This key allows to set the language for a structure element. The value should be a bcp-identifier, e.g. `|de-De|`.

ref This key allows to add references to other structure elements, it adds the `|/Ref|` array to the structure. The value should be a comma separated list of structure labels set with the `|label|` key. e.g. `|ref=label1,label2|`.

E This key sets the `|/E|` key, the expanded form of an abbreviation or an acronym (I couldn't think of a better name, so I stuck to E).

```
\tagstructend
\tag_struct_end:
```

These commands end a structure. They don't end a group and it doesn't matter if they are in another group as the starting commands.

```
\tagstructuse{<label>}
\tag_struct_use:n{<label>}
```

These commands insert a structure previously stashed away as `kid` into the currently active structure. A structure should be used only once, if the structure already has a parent you will get a warning.

4.3.4. Root structure

A document should have at least one structure which contains the whole document. A suitable tag is `Document` or `Article`. I'm considering to automatically inserting it.

4.3.5. Attributes and attribute classes

Structure Element can have so-called attributes. A single attribute is a dictionary (or a stream but this is currently not supported by the package as I don't know an use-case) with at least the required key `/Owner` which describes the scope the attribute applies too. As an example here an attribute that can be attached to tabular header (type TH) and adds the info that the header is a column header:

```
<</Owner /Table /Scope /Column>>
```

One or more such attributes can be attached to a structure element. It is also possible to store such an attribute under a symbolic name in a so-called “ClassedMap” and then to attach references to such classes to a structure.

To use such attributes you must at first declare it in `\tagpdfsetup` with the key `newattribute`. This key takes two argument, a name and the content of the attribute. The name should be a sensible key name, the content a dictionary.

```
\tagpdfsetup
{
  newattribute =
    {TH-col}{/Owner /Table /Scope /Column},
  newattribute =
    {TH-row}{/Owner /Table /Scope /Row},
}
```

Attributes are only written to the PDF when used, so it is not a problem to predeclare a number of standard attributes.

It is your responsibility that the content of the dictionary is valid PDF and that the values are sensible!

Attributes can then be used with the key `attribute` or `attribute-class` which both take a comma list of attribute names as argument:

```
\tagstructbegin{tag=TH,
  attribute-class= {TH-row,TH-col},
  attribute = {TH-row,TH-col},
}
```

4.4. Task 3: tree Management

When all the document content has been correctly marked and the data for the trees has been collected they must be flushed to the PDF. This is done automatically (if the package has been activated) with the following command in `\AfterEndDocument`:

`\tag_finish_structure:`

This will hopefully write all the needed objects and values to the PDF. (Beside the already mentioned StructTreeRoot and StructElem objects, additionally a so-called ParentTree is needed which records the parents of all the marked contents bits, a Rolemap, perhaps a ClassMap and object for the attributes, and a few more values and dictionaries).

I'm not quite sure if this shouldn't be a really internal command.

4.5. A fully marked up document body

The following shows the marking needed for a section, a sentence and a list with two items. It is obvious that one wouldn't like to have to do this for real documents. If tagging should be usable, the commands must be hidden as much as possible inside suitable \TeX commands and environments.

```
\begin{document}

\tagstructbegin{tag=Document}

\tagstructbegin{tag=Sect}
\tagstructbegin{tag=H}
\tagmcbegin{tag=H} %avoid page break!
\section{Section}
\tagmccend
\tagstructend
\tagstructbegin{tag=P}
\tagmcbegin{tag=P,raw=/Alt (x)}
a paragraph\par x
\tagmccend
\tagstructend

\tagstructbegin{tag=L} %List
\tagstructbegin{tag=LI}
\tagstructbegin{tag=Lbl}
\tagmcbegin{tag=Lbl}
1.
\tagmccend
\tagstructend
\tagstructbegin{tag=LBody}
\tagmcbegin{tag=P}
List item body
\tagmccend
\tagstructend %lbody
\tagstructend %Li
```



```

\tagstructbegin{tag=LI}
\tagstructbegin{tag=Lbl}
\tagmcbegin{tag=Lbl}
2.
\tagmcend
\tagstructend
\tagstructbegin{tag=LBody}
\tagmcbegin{tag=P}
another List item body
\tagmcend
\tagstructend %lbody
\tagstructend %Li
\tagstructend %L

\tagstructend %Sect
\tagstructend %Document
\end{document}

```

4.6. Lazy and automatic tagging

A number of features of PDF readers need a fully tagged PDF. As an example screen readers tend to ignore alternative text (see section 5) if the PDF is not fully tagged. Also reflowing a PDF only works for me (even if real space chars are in the PDF) if the PDF is fully tagged.

This means that even if you don't care about a proper structure you should try to add at least some minimal tagging. With pdf_latex this is not easy due to the page break problem. But with lua_latex you can use an Document structure and inside it rather large mc-chunks. This minimizes the needed work.

One could ask if in lua mode the code couldn't try to mark up unmarked parts e.g. as P-type chunks, like it marks them up as artifacts currently. Sadly this is not so easy, as it is quite difficult to reliably identify the structure and the place in the kids array where such chunks belongs too. I also don't think that it is really needed. It is not so difficult to define user macros which e.g. opens a structure and start an mc-chunk or which close an open mc-chunk before issuing the next `\tagmcbegin`.

4.7. Adding tagging to commands

As mentioned above the mc-markers should not be nested. Basically you write:

```

\tagmcbegin{..}some text ... \tagmcend
<optional structure commands>
\tagmcbegin{..}some other text \tagmcend

```

This is quite workable as long as you mark everything manually. But how to write commands, e.g. for a tabular or a graphic, that do tagging automatically without breaking the flow and the structure?

5. Alternative text, ActualText and text-to-speech software

The PDF format allows to add alternative text through the `/Alt` and the `/ActualText` key. There is also the `/E` key for acronyms but I will ignore it for now. Both can be added either to the marked content in the page stream or to the object describing the structure.

The value of `/ActualText` (inserted by tagpdf with `actualtext`) is meant to replace single characters or rather small pieces of text. It can be used also without any tagging (e.g. with the package `accsupp`). If the PDF reader support this (adobe reader does, sumatra not) one can change with it how a piece of text is copied and pasted e.g. to split up a ligature.

`/Alt` (inserted by tagpdf with `alttext`) is a key to improve accessibility: with it one can add to a picture or something else an alternative text.

The file `ex-alt-actualtext.tex` shows some experiments I made with with both keys and text-to-speech software (the in-built of adobe and nvda). To sum them up:

- The keys have an impact on text-to-speech software only if the document is fully tagged.
- `/ActualText` should be at best used around short pieces of marked content.
- `/Alt` is used at best with a structure – this avoids problems with luatex where marked contents blocks can be split over pages.
- To some extend one can get a not so bad reading of math with the alternative text.

6. Standard type and new tags

The PDF reference mentions a number of standard types: Document, Part, Art, Sect, Div, BlockQuote, Caption, TOC, TOCI, Index, NonStruct, H, H1, H2, H3, H4, H5, H6, P, L, LI, Lbl, LBody, Table, TR, TH, TD, THead, TBody, TFoot, Span, Quote, Note, Reference, BibEntry, Code, Link, Annot, Figure, Formula, Form, Artifact

Their meaning can be looked up in the PDF-reference [1].

New tags can be defined in the setup command with the key `add-new-tag`. It takes a value consisting of two names separated by a slash. The first is the new name, the second a known (e.g. a standard) tag it should be mapped too. Example:

```
\tagpdfsetup{add-new-tag = section/H1}
```

7. “Real” space glyphs

TeX uses only spaces (horizontal movements) to separate words. That means that a PDF reader has to use some heuristic when copying text or reflowing the text to decide if a space is meant as a word boundary or e.g. as a kerning. Accessible document should use real space glyphs (U+0032) from a font in such places.

With the key `interwordspace` you can activate such space glyphs.

With pdfTeX this will simply call the primitive `\pdfinterwordspaceon`. pdfTeX will then insert at various places a char from a font called dummy-space. Attention! This means that at every space there are additional font switches in the PDF: from the current font to the dummy-space font and back again. This will make the PDF larger. As `\pdfinterwordspaceon` is a primitive function it can't be fine tuned or adapted. You can only turn it on and off and insert manually such a space glyph with `\pdffakepace`.

With LuaTeX (in Lua mode) `interwordspace` is implemented with a Lua-function which is inserted in two callbacks and marks up the places where it seems sensible to insert a space glyph. Later in the process (when also the mc-markers are inserted) the space glyphs are injected – the code will take the glyph from the current font if this has a space glyph or switch to the default Latin Modern font. The current code works reasonable well in normal text.

The key `show-spaces` will show lines at the places where in Lua mode spaces are inserted and so can help you to find problematic places. For listings – which have a quite specific handling of spaces – you can find a suggestion in the example `ex-space-glyph-listings`.

Attention: Even with real spaces copy& pasting of code doesn't need to give the correct results: you get spaces but not necessarily the right number of spaces. The PDF viewers I tried all copied four real space glyphs as one space. I only got the four spaces with the export to text or XML in the AdobePro.

`\pdffakepace`

This is in pdfTeX a primitive. It inserts the dummy space glyph. `tagpdf` defines this command also for LuaTeX – attention if it can perhaps insert break points.

8. Accessibility is not only tagging

A tagged PDF is needed for accessibility but this is not enough. As already mentioned there are more requirements:

- The language must be declared by adding a `/Lang xx-XX` to the PDF catalog or – if the language changes for a part of the text to the structure or the marked content. Setting the document language can be rather easily done with existing packages. With the new PDF resource management it should be done with `\pdfmanagement_add:nnn{Catalog}{Lang}{(en-US)}`. For settings in marked content and structure I will have to add keys.

- All characters must have an unicode representation or a suitable alternative text. With lualatex and open type (unicode) fonts this is normally not a problem. With pdflatex it could need

```
\input{glyphtounicode}
\pdfgentounicode=1
```

and perhaps some\pdfglyphtounicode commands.

- Hard and soft hyphen must be distinct.
- Spaces between words should be space glyphs and not only a horizontal movement. See section 7.
- Various small infos must be present in the catalog dictionary, info dictionary and the page dictionaries, e.g. metadata like title.

If suitable I will add code for this tasks to this packages. But some of them can also be done already with existing packages like hyperref, hyperxmp, pdfx.

9. To-do

- Add commands and keys to enable/disable the checks.
- Check/extend the code for language tags.
- Think about math.
- Think about Links/Annotations
- Keys for alternative and actualtext. How to define the input encoding? Like in Accsupp?
- Check twocolumn documents
- Examples
- Write more Tests
- Write more Tests
- Unicode
- Hyphenation char
- Think about included (tagged) PDF. Can one handle them?
- Improve the documentation
- Tag as proof of concept the documentation (nearly done)
- Document the code better
- Create dtx
- Find someone to check and improve the lua code

- Move more things to lua in the luamode
- Find someone to check and improve the rest of the code
- Check differences between PDF versions 1.7 and 2.0.
- bidi?

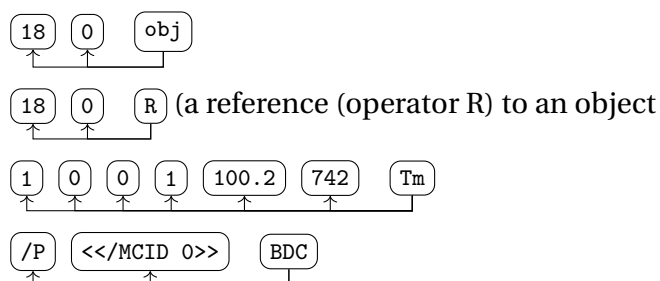
References

- [1] Adobe Systems Incorporated. *PDF Reference, sixth edition*. 2006. URL: https://www.adobe.com/content/dam/acom/en/devnet/pdf/pdfs/PDF32000_2008.pdf.
- [2] TeX User Group. *PDF accessibility and PDF standards*. URL: <https://tug.org/twg/accessibility/>.
- [3] veraPDF consortium. *veraPDF*. URL: <http://verapdf.org/>.
- [4] Zugang für alle – Schweizerische Stiftung zur behindertengerechten Technologienutzung. *PDF Accessibility Checker (PAC 3)*. URL: <http://www.access-for-all.ch/ch/pdf-werkstatt/pdf-accessibility-checker-pac.html> (visited on 07/05/2018).

A. Some remarks about the PDF syntax

This is not meant as a full reference only as a background to make the examples and remarks easier to understand.

postfix notation PDF uses in various places postfix notation. This means that the operator is behind its arguments:



Names PDF knows a sort of variable called a “name”. Names start with a slash and may include any regular characters, but not delimiter or white-space characters. Uppercase and lowercase letters are considered distinct: /A and /a are different names. /.notdef and /Adobe#20Green are valid names.

Quite a number of the options of `tagpdf` actually define such a name which is later added to the PDF. I recommend *strongly* not to use spaces and exotic chars in such names. While it is possible to escape such names it is rather a pain when moving them through the various lists and commands and quite probably I forgot some place where it is needed.

Strings There are two types of strings: *Literal strings* are enclosed in round parentheses. They normally contain a mix of ascii chars and octal numbers:

```
(gr\374\377ehello[]\050\051).
```

Hexadezimal strings are enclosed in angle brackets. They allow for a representation of all characters the whole unicode ranges. This is the default output of lualatex.

```
<003B00600243013D0032>.
```

Arrays Arrays are enclosed by square brackets. They can contain all sort of objects including more arrays. As an example here an array which contains five objects: a number, an object reference, a string, a dictionary and another array. Be aware that despite the spaces `15 0 R` is *one* element of the array.

```
[0 15 0 R (hello) <</Type /X>> [1 2 3]]
```

```
(0) (15 0 R) (hello) (<</Type /X>>) ([1 2 3])
```

Dictionaries Dictionaries are enclosed by double angle brackets. They contain key-value pairs. The key is always a name. The value can be all sort of objects including more dictionaries. It doesn't matter in which order the keys are given.

Dictionaries can be written all in one line:

```
<</Type/Page/Contents 3 0 R/Resources 1 0 R/Parent 5 0 R>> but at least for ex-  
amples a layout with line breaks and indentation is more readable:
```

```
<<  
  /Type      /Page  
  /Contents  3 0 R  
  /Resources 1 0 R  
  /MediaBox  [0 0 595.276 841.89]  
  /Parent    5 0 R  
>>
```

(indirect) objects These are enclosed by the keywords `obj` (which has two numbers as prefix arguments) and `endobj`. The first argument is the object number, the second a generation number – if a PDF is edited objects with a larger generation number can be added. As with `pdflatex/lualatex` the PDF is always new we can safely assume that the number is always 0. Objects can be referenced in other places with the `R` operator. The content of an object can be all sort of things.

streams A stream is a sequence of bytes. It can be long and is used for the real content of PDF: text, fonts, content of graphics. A stream starts with a dictionary which at least sets the `/Length` name to the length of the stream followed by the stream content enclosed by the keywords `stream` and `endstream`.

Here an example of a stream, an object definition and reference. In the object 2 (a page object) the `/Contents` key references the object 3 and this then contains the text of the page in a stream. `Tf`, `Tm` and `TJ` are (postfix) operators, the first chooses the font with

the name /F15 at the size 10.9, the second displaces the reference point on the page and the third inserts the text.

```
% a page object (shortened)
2 0 obj
<<
  /Type/Page
  /Contents 3 0 R
  /Resources 1 0 R
  ...
>>
endobj

%the /Contents object (/Length value is wrong)
3 0 obj
<</Length 153 >>
stream
BT
  /F15 10.9 Tf 1 0 0 1 100.2 746.742 Tm [(hello)]TJ
ET
endstream
endobj
```

In such a stream the BT–ET pair encloses texts while drawing and graphics are outside of such pairs.

Number tree This is a more complex data structure that is meant to index objects by numbers. In the core is an array with number-value pairs. A simple version of number tree which has the keys 0 and 3 is

```
6 0 obj
<<
  /Nums [
    0 [ 20 0 R 22 0 R]
    3 21 0 R
  ]
>>
endobj
```

This maps 0 to an array and 2 to the object reference 21 0 R. Number trees can be split over various nodes – root, intermediate and leaf nodes. We will need such a tree for the *parent tree*.