

Context

Musings

hans hagen

Content

Introduction	4
1 Children of \TeX	6
2 Advertising \TeX	32
3 Why use \TeX ?	36
4 What's to stay, what's to go	44
5 Stability	52
6 Meta \TeX , a roadmap	60
7 What's in a name	68
8 About what Con \TeX t isn't	74
9 False promises	84
10 About manuals	92
11 Performance again	96
12 All those \TeX 's	104
13 Hidden treasures	110
14 Don't use \TeX !	114
15 Speeding up \TeX	120
16 Unicode	132
17 Con \TeX t in \TeX Live 2023	160
18 How not to install Con \TeX t	166
19 Side effects	168
20 Strange assumptions	176
21 It's about time	180

22	Missing kerns	192
23	We deserved it	200
24	Testing RiscV	206
25	(Ab)use of font features	210
26	The shift (or: just moving on)	220
27	Illiterate programming	224
28	Fake quality	230

Introduction

This is a collection of articles and wrap-ups that don't suit in other manuals or collections. Some are published, some meant as draft for a presentation.

The "Children of \TeX " article is the framework for a presentation at Bacho \TeX 2017 in Poland, and covers the main theme of the conference. In the aftermath of that conference I wrote "Advertising \TeX " and later "Why use \TeX ?". The 2018 Bacho \TeX conference theme is explored in "What's to stay, what's to go". After a short discussion on the Con \TeX t mailing list about stability (at the moment that MkII had been frozen for more than a decade but is still used without problems) I wrote "Stability".

Many of the thoughts in these articles are influenced by discussions with my colleagues Ton Otten and Kees van Marle, users and developers. Operating in a similar arena, they provide me the reflection needed to sort out my thoughts on these matters.

The order in this document is not chronological. In the meantime we also put some development related stories in this collection, just because they have to fit in somewhere.

Not all musings are checked and copy-edited so let me know if there are errors and typos in them.

Hans Hagen
Hasselt NL
2017–2023+

Children of T_EX 1

1.1 The theme

Nearly always T_EX conferences carry a theme. As there have been many conferences the organizers have run out of themes involving fonts, macros and typesetting and are now cooking up more fuzzy ones. Take the Bachotug 2017 theme:

Premises The starting point, what we have, what do we use, what has been achieved?

Predilections How do we act now, how do we want to act, what is important to us and what do we miss?

Predictions What is the future of T_EX, what we'd like to achieve and can we influence it?

My first impression with these three P words was: what do they mean? Followed by the thought: this is no longer a place to take kids to. But the Internet gives access to the Cambridge Dictionary, so instead of running to the dusty meter of dictionaries somewhere else in my place, I made sure that I googled the most recent definitions:

premise an idea or theory on which a statement or action is based

predilection if someone has a predilection for something, they like it a lot

prediction a statement about what you think will happen in the future

I won't try to relate these two sets of definitions but several words stand out in the second set: idea, theory, action, like, statement and future. Now, as a preparation for the usual sobering thoughts that Jerzy, Volker and I have when we staring into a BachoT_EX campfire I decided to wrap up some ideas around these themes and words. The books that I will mention are just a selection of what you can find distributed around my place. This is not some systematic research but just the result of a few weeks making a couple of notes while pondering about this conference.

1.2 Introduction

One cannot write the amount of T_EX macros that I've written without also liking books. If you look at my bookshelves the topics are somewhat spread over the possible spectrum of topics: history, biology, astronomy, paleontology, general science but surprisingly little math. There are a bunch of typography-related books but only some have been read: it's the visuals that matter most and as there are no real developments I haven't bought new ones in over a decade, although I do buy books that look nice for our office display but the content should be interesting too. Of course I do have a couple of books about computer (related) science and technology but only a few are worth a second look. Sometimes I bought computer books expecting to use them (in some

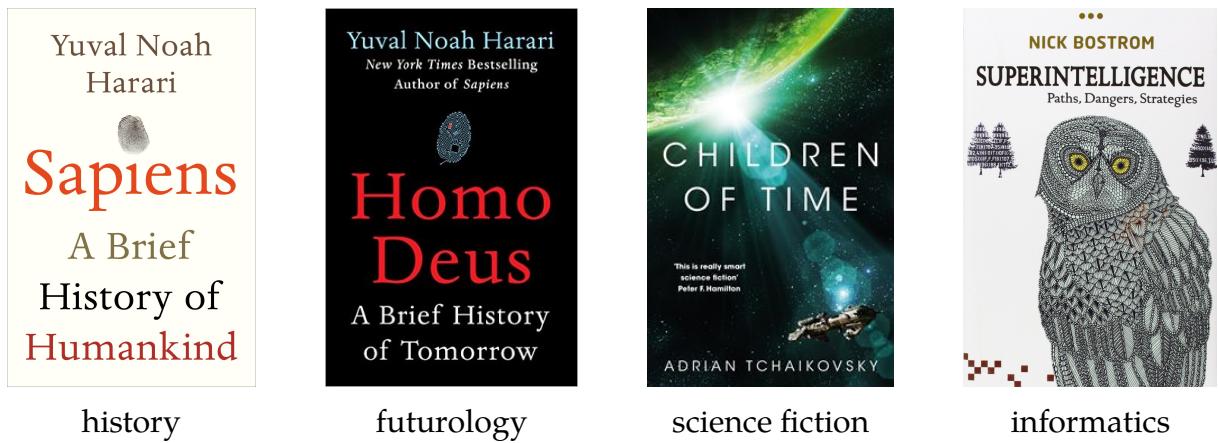


Figure 1.1

project) but I must admit that most have not been read and many will soon end up in the paper bin (some already went that way). I'll make an exception for Knuth, Wirth and a few other fundamental ones that I (want to) read. And, I need to catch up on deep learning, so that might need a book.

My colleagues and I have many discussions, especially about what we read, and after a few decades one starts seeing patterns. Therefore the last few years it was a pleasant surprise for me to run into books and lectures that nicely summarize what one has noticed and discussed in a consistent way. My memory is not that good, but good enough to let some bells ring.

The first book that gave me this “finally a perfect summary of historic developments” feeling is “Sapiens” by Yuval Noah Harari. The author summarizes human history from a broad perspective where modern views on psychology, anthropology and technical developments are integrated. It’s a follow up on a history writing trend started by Jared Diamond. The follow up “Homo Deus” looks ahead and is just as well written. It also integrates ideas from other fields, for instance those related to development of artificial intelligence (Dennett, Bostrom, etc.).

Another inspiration for this talk and article is the 50 hour lecture series on behavioral biology by Robert Sapolsky of Stanford University, brought to my attention by my nephew Bram who visited a few \TeX conferences with me and who is now also forced to use \TeX for assignments and reports. (How come self-published books used at universities often look so bad?)

The title of this talk is inspired by the book “Children of Time” by Adrian Tchaikovsky that I read recently. There are science fiction writers who focus on long term science and technology, such as some of Alastair Reynolds, while others follow up on recent development in all kind of sciences. One can recognize aspects of “Superintelligence” by Bostrom in Neal Asher’s books, insights in psychology in the older Greg Bear books, while in the mentioned “Children of Time” (socio)biological insights dominate. The main thread in that book is the development of intelligence, social behaviour, language, script and cooperation in a species quite different from us: spiders. It definitely avoids the anthropocentric focus that we normally have.

So how does this relate to the themes of the BachtEX conference? I will pick out some ways to approach them using ideas from the kind of resources mentioned above. I could probably go on and on for pages because once you start relating what you read and hear to this TeX ecosystem and community, there is no end. So, consider this a snapshot, that somehow relates to the themes:

premise Let's look at what the live sciences have to say about TeX and friends and let's hope that I don't offend the reader and the field.

predilection Let's figure out what brings us here to this place deeply hidden in the woods, a secret gathering of the TeX sect.

prediction Let's see if the brains present here can predict the future because after all, according to Dennett, that is what brains are for.

At school I was already intrigued by patterns in history: a cyclic, spiral and sinusoid social evolution instead of a pure linear sequence of events. It became my first typeset-by-typewriter document: Is history an exact science? Next I will use and abuse patterns and ideas to describe the TeX world, not wearing a layman's mathematical glasses, but more from the perspective of live sciences, where chaos dominates.

1.3 The larger picture

History of mankind can be roughly summarized as follows. For a really long time we were hunters but at some point (10K years ago) became farmers. As a result we could live in larger groups and still feed them. The growing complexity of society triggered rules and religion as instruments for stability and organization (I use the term religion in its broadest sense here). For quite a while cultures came and went, and climate changes are among the reasons.

After the industrial revolution new religions were invented (social, economic and national liberalism) and we're now getting dataism (search for Harari on youtube for a better summary). Some pretty great minds seem to agree that we're heading to a time when humans as we are will be outdated. Massive automation, interaction between the self and computer driven ecosystems, lack of jobs and purpose, messing around with our genome. Some countries and cultures still have to catch up on the industrial revolution, if they manage at all, and maybe we ourselves will be just as behind reality soon. Just ask yourself: did you manage to catch up? Is TeX a stone age tool or a revolutionary turning point?

A few decades ago a trip to BachtEX took more than a day. Now you drive there in just over half a day. There was a time that it took weeks: preparation, changing horses, avoiding bad roads. Not only your own man-hours were involved. It became easier later (my first trip took only 24 hours) and recently it turned into a piece of cake: you don't pick up maps but start your device; you don't need a travel agent but use the Internet;

there are no border patrols, you can just drive on. (Okay, maybe some day soon border patrols at the Polish border show up again, just like road tax police in Germany, but that might be a temporary glitch.)

Life gets easier and jobs get lost. Taxi and truck drivers, travel agents, and cashiers become as obsolete as agricultural workers before. Next in line are doctors, lawyers, typesetters, printers, and all those who think they're safe. Well, how many people were needed 400 years ago to produce the proceedings of a conference like this in a few days' time span? Why read the introduction of a book or a review when you can just listen to the author's summary on the web? How many conferences still make proceedings (or go for videos instead), will we actually need editors and typesetters in the future? How much easier has it become to design a font, including variants? What stories can designers tell in the future when programs do the lot? The narrower your speciality is, the worse are your changes; hopefully the people present at this conference operate on a broader spectrum. It's a snapshot. I will show some book covers as reference but am aware that years ago or ahead the selection could have been different.

1.4 Words

Words (whatever they represent) found a perfect spot to survive: our minds. Then they made it from speech (and imagination) into writing: carved in stone, wood, lead. At some point they managed to travel over wires but no matter what happened, they are still around. Typesetting as visualization is also still surrounding us so that might give us a starting point for ensuring a future for \TeX to work on, because \TeX is all about words. There is a lot we don't see; imagine if our eyes had microscopic qualities. What if we could hear beyond 20KHz. Imagine we could see infrared. How is that with words. What tools, similar in impact as \TeX , can evolve once we figure that out. What if we get access to the areas of our brain that hold information? We went from print to screen and \TeX could cope with that. Can it cope with what comes next?

The first printing press replaced literal copying by hand. Later we got these linotype-like machines but apart from a few left, these are already thrown out of windows (as we saw in a movie a few Bacho \TeX 's ago). Photo-typesetting has been replaced too and because a traditional centuries old printing press is a nice to see item, these probably ring more bells than that gray metal closed box typesetters. Organizers of \TeX conferences love to bring the audience to old printing workshops and museums. At some point computers got used for typesetting and in that arena \TeX found its place. These gray closed boxes are way less interesting than something mechanical that at least invites us to touch it. How excited can one be about a stack of \TeX Live dvds?

1.5 Remembering

Two times I visited the part of the science museum in London with young family members: distracted by constantly swiping their small powerful devices, they didn't have the least interest in the exhibited computer related items, let alone the fact that the couch

they were sitting on was a Cray mainframe. Later on, climbing on some old monument or an old cannon seemed more fun. So, in a few decades folks will still look at wooden printing presses but quickly walk through the part of an exhibition where the tools that we use are shown. We need to find ways to look interesting. But don't think we're unique: how many kids find graphical trend-setting games like *Myst* and *Riven* still interesting? On the other hand a couple of month ago a bunch of nieces and nephews had a lot of fun with an old Atari console running low-res bitmap games. Maybe there is hope for good old \TeX .

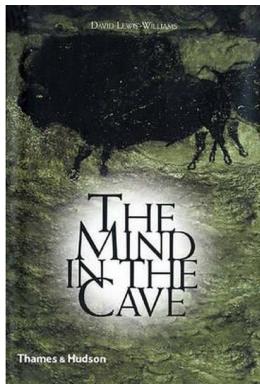
If indeed we're heading to a radically different society one can argue if this whole discussion makes sense. When the steam engine showed up, the metaphor for what went on in our heads was that technology, It's a popular example of speakers on this topic: "venting off steam". When electricity and radio came around metaphors like "being on the same wavelength" showed up. A few decades ago the computer replaced that model although in the meantime the model is more neurobiological: we're a hormone and neurotransmitter driven computer. We don't have memory the way computers do.

How relevant will page breaks, paragraph and line breaks be in the future? Just like "venting off steam" may make no sense to the youth, asking a typesetter to "give me a break" might not make much sense soon. However, when discussing automated typesetting the question "are we on the same page" still has relevance.

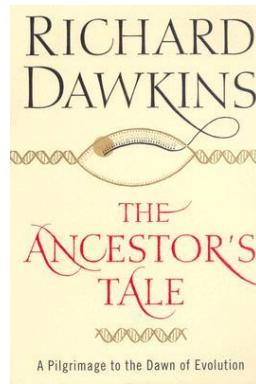
Typesetting with a computer might seem like the ultimate solution but it's actually rather dumb when we consider truly intelligent systems. On the large scale of history and developments what we do might get quite unnoticed. Say that mankind survives the next few hundred years one way or the other. Science fiction novels by Jack McDevitt have an interesting perspective of rather normal humans millennia ahead of us who look back on these times in the same way as we look back now. Nothing fundamental changed in the way we run society. Nearly nothing from the past is left over and apart from being ruled by ais people still do sort of what they do now. \TeX ? What is that? Well, there once was this great computer scientist Knuth (in the remembered row of names like Aristotle —I just started reading "The Lagoon" by Armand Leroi— Newton, Einstein, his will show up) who had a group of followers that used a program that he seems to have written. And even that is unlikely to be remembered, unless maybe user groups manage to organize an archive and pass that on. Maybe the fact that \TeX was one of the first large scale open source programs, of which someone can study the history, makes it a survivor. The first program that was properly documented in detail! But then we need to make sure that it gets known and persists.

1.6 Automation

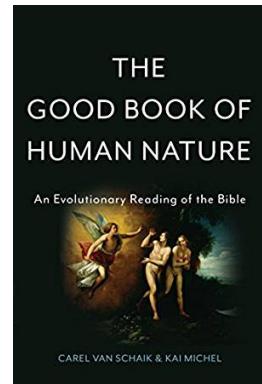
In a recent interview Daniel Dennett explains that his view of the mind as a big neural network, one that can be simulated in software on silicon, is a bit too simplistic. He wonders if we shouldn't more tend to think of a network of (selfish) neurons that group together in tasks and then compete with each other, if only because they want to have something to do.



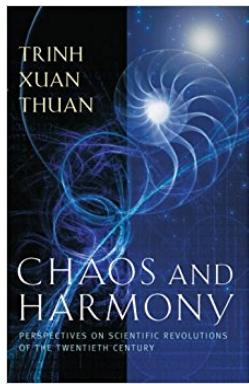
paleontology



evolutionary
biology



anthropology



physics

Figure 1.2

Maybe attempts to catch the creative mindset and working of a typesetter in algorithms is futile. What actually is great typography or good typesetting? Recently I took a look at my bookshelf wondering what to get rid of — better do that now than when I'm too old to carry the crap down (crap being defined as uninteresting content or bad looking). I was surprised about the on-the-average bad quality of the typesetting and print. It's also not really getting better. One just gets accustomed to what is the norm at a certain point. Whenever they change the layout and look and feel of the newspaper I read the arguments are readability and ease of access. Well, I never had such a hard time reading my paper as today (with my old eyes).

Are we, like Dennett, willing to discard old views on our tools and models? When my first computer was a RCA 1802 based kit, that had 256 bytes of memory. My current laptop (from 2013) is a Dell Precision workstation with an extreme quad core processor and 16 GB of memory and ssd storage. Before I arrived there I worked with dec-10, vax and the whole range of Intel cpus. So if you really want to compare a brain with a computer, take your choice.

I started with \TeX on a 4 MHz desk top with 640 MB memory and a 10 MB hard disk. Running Con \TeX Xt MkIV with Lua \TeX on such a machine is no option at all, but I still carry the burden of trying to write efficient code (which is still somewhat reflected in the code that makes up Con \TeX Xt). In the decades that we have been using \TeX we had to adapt! Demands changed, possibilities changed, technologies changed. And they keep changing. How many successive changes can a \TeX user handle? Sometimes, when I look and listen I wonder.

If you look back, that is, if you read about the tens of thousands of years that it took humans to evolve ("The mind in the cave" by Lewis-Williams is a good exercise) you realize even more in what a fast-paced time we live and that we're witnessing transitions of another magnitude.

In the evolution of species some tools were invented multiple times, like eyes. You see the same in our \TeX world: multiple (sub)macro packages, different font technologies, the same solutions but with an alternative approach. Some disappear, some

stay around. Just like different circumstances demand different solutions in nature, so do different situations in typesetting, for instance different table rendering solutions. Sometime I get the feeling that we focus too much on getting rid of all but one solution while more natural would be to accept diversity, like bio-diversity is accepted. Transitions nowadays happen faster but the question is if, like aeons before, we (have to) let them fade away. When evolution is discussed the terms ‘random’, ‘selection’, ‘fit’, and so on are used. This probably also applies to typography: at some point a font can be used a lot, but in the end the best readable and most attractive one will survive. Newspapers are printed in many copies, but rare beautiful books hold value. Of course, just like in nature some developments force the further path of development, we don’t suddenly grow more legs or digits on our hands. The same happens with \TeX on a smaller timescale: successors still have the same core technology, also because if we’d drop it, it would be something different and then give a reason to reconsider using such technology (which likely would result in going by another path).

1.7 Quality

Richard Dawkins “The Ancestor’s Tale” is a non-stop read. In a discussion with Jared Diamond about religion and evolution they ponder this thread: you holding the hand of your mother who is handing her mother’s hand and so on till at some point fish get into the picture. The question then is, when do we start calling something human? And a related question is, when does what we call morality creeps in? Is 50% neanderthal human or not?

So, in the history of putting thoughts on paper: where does \TeX fit in? When do we start calling something automated typesetting? When do we decide that we have quality? Is \TeX so much different from its predecessors? And when we see aspects of \TeX (or related font technology) in more modern programs, do we see points where we cross qualitative or other boundaries? Is a program doing a better job than a human? Where do we stand? There are fields where there is no doubt that machines outperform humans. It’s probably a bit more difficult in aesthetic fields except perhaps when we lower the conditions and expectations (something that happens a lot).

For sure \TeX will become obsolete, maybe even faster than we think, but so will other typesetting technologies. Just look back and have no illusions. Till then we can have our fun and eventually, when we have more free time than we need, we might use it out of hobbyism. Maybe \TeX will be remembered by probably its most important side effect: the first large scale open source, the time when users met over programs, Knuth’s disciples gathered in user groups, etc. The tools that we use are just a step in an evolution. And, as with evolution, most branches are pruned. So, when in the far future one looks back, will they even notice \TeX ? The ancestor’s tail turns the tree upside down: at the end of the successful branch one doesn’t see the dead ends.

Just a thought: cds and media servers are recently being replaced (or at least accompanied) by Long Play records. In the shop where I buy my cds the space allocated to

records grows at the cost of more modern media. So, maybe at some point retro-type-setting will pop up. Of course it might skip \TeX and end up at woodcutting or printing with lead.

1.8 What mission

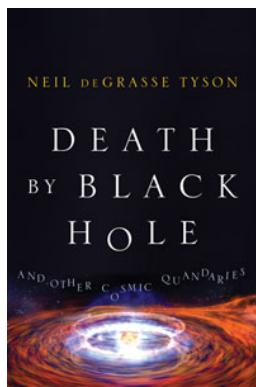
We rely on search engines instead of asking around or browsing libraries. Do students really still read books and manuals or do they just search and listen to lectures. Harari claims that instead of teaching kids facts in school we should just take for granted that they can get all the data they want and that we should learn them how to deal with data and adapt to what is coming. We take for granted that small devices with human voices show us the route to drive to Bacho \TeX , for instance, although by now I can drive it without help. In fact, kids can surprise you by asking if we're driving in Germany when we are already in Poland.

We accept that computer programs help physicians in analyzing pictures. Some wear watches that warn them about health issues, and I know a few people who monitor their sugar levels electronically instead of relying on their own measurements. We seem to believe and trust the programs. And indeed, we also believe that \TeX does the job in the best way possible. How many people really understand the way \TeX works?

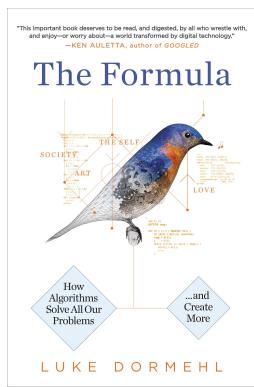
We still have mailing lists where we help each other. There are also wikis and forums like stack exchange. But who says that even a moderate bit of artificial intelligence doesn't answer questions better. Of course there needs to be input (manuals, previous answers, etc.) but just like we need fewer people as workforce soon, the number of experts needed also can be smaller. And we're still talking about a traditional system like \TeX . Maybe the social experience that we have on these media will survive somehow, although: how many people are members of societies, participate in demonstrations, meet weekly in places where ideas get exchanged, compared to a few decades ago? That being said, I love to watch posts with beautiful Con $\text{\TeX}t$ solutions or listen to talks by enthusiastic users who do things I hadn't expected. I really hope that this property survives, just like I hope that we will be able to see the difference between a real user's response and one from an intelligent machine (an unrealistic hope I fear). Satisfaction wins and just like our neurological subsystems at some point permanently adapt to thresholds (given that you trigger things often enough), we get accustomed to what \TeX provides and so we stick to it.

1.9 Intelligence versus consciousness

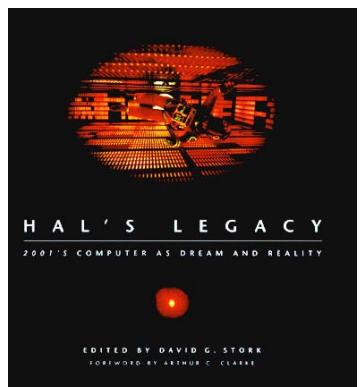
Much of what we do is automated. You don't need to think of which leg to move and what foot to put down when you walk. Reacting to danger also to a large extent is automated. It doesn't help much to start thinking about how dangerous a lion can be when it's coming after you, you'd better move fast. Our limbic system is responsible for such automated behaviour, for instance driven by emotions. The more difficult tasks and thoughts about them happen in the frontal cortex (sort of).



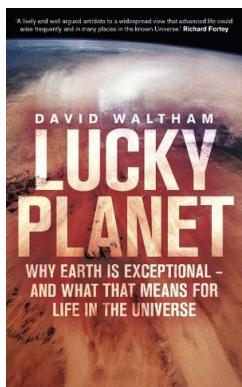
astronomy



informatics



future science



earth science

Figure 1.3

For most users \TeX is like the limbic system: there is not much thinking involved, and the easy solutions are the ones used. Just like hitting a nerve triggers a chain of reactions, hitting a key eventually produces a typeset document. Often this is best because the job needs to get done and no one really cares how it looks; just copy a preamble, key in the text and assume that it works out well (enough). It is tempting to compare \TeX 's penalties, badness and other parameters with levels of hormones and neurotransmitters. Their function depends on where they get used and the impact can be accumulated, blocked or absent. It's all magic, especially when things interact.

Existing \TeX users, developers and user groups of course prefer to think otherwise, that it is a positive choice by free will. That new users have looked around and arrived at \TeX for good reason: their frontal cortex steering a deliberate choice. Well, it might have played a role but the decision to use \TeX might in the end be due to survival skills: I want to pass this exam and therefore I will use that weird system called \TeX .

All animals, us included, have some level of intelligence but also have this hard to describe property that we think makes us what we are. Intelligence and consciousness are not the same (at least we know a bit about the first but nearly nothing about the second). We can argue about how well composed some music is but why we like it is a different matter.

We can make a well thought out choice for using \TeX for certain tasks but can we say why we started liking it (or not)? Why it gives us pleasure or maybe grief? Has it become a drug that we got addicted to? So, one can make an intelligent decision about using \TeX but getting a grip on why we like it can be hard. Do we enjoy the first time struggle? Probably not. Do we like the folks involved? Yes, Don Knuth is a special and very nice person. Can we find help and run into a friendly community? Yes, and a unique one too, annoying at times, often stimulating and on the average friendly for all the odd cases running around.

Artificial intelligence is pretty ambitious, so speaking of machine intelligence is probably better. Is \TeX an intelligent program? There is definitely some intelligence built in and the designer of that program is for sure very intelligent. The designer is also a conscious entity: he likes what he did and finds pleasure in using it. The program on

the other hand is just doing its job: it doesn't care how it's done and how long it takes: a mindless entity. So here is a question: do we really want a more intelligent program doing the job for us, or do those who attend conferences like BachoTeX enjoy TeXing so much that they happily stay with what they have now? Compared to rockets tumbling down and/or exploding or Mars landers thrashing themselves due to programming errors of interactions, TeX is surprisingly stable and bug free.

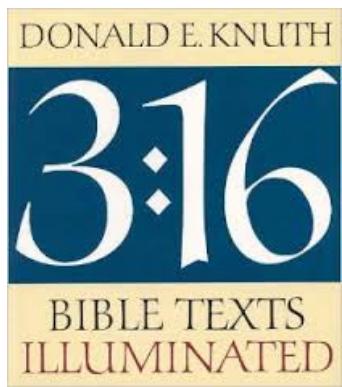
1.10 Individual versus group evolution

After listening for hours to Sapolsky you start getting accustomed to remarks about (unconscious) behaviour driven by genes, expression and environment, aimed at "spreading many copies of your genes". In most cases that is an individual's driving force. However, cooperation between individuals plays a role in this. A possible view is that we have now reached a state where survival is more dependent on a group than on an individual. This makes sense when we consider that developments (around us) can go way faster than regular evolution (adaptation) can handle. We take control over evolution, a mechanism that needs time to adapt and time is something we don't give it anymore.

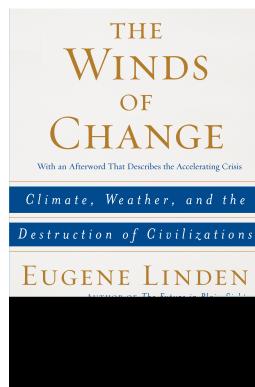
Why does TeX stay around? It started with an individual but eventually it's the groups that keeps it going. A too-small group won't work but too-large groups won't work either. It's a known fact that one can only handle some 150 social contacts: we evolved in small bands that split when they became too large. Larger groups demanded abstract beliefs and systems to deal with the numbers: housing, food production, protection. The TeX user groups also provide some organization: they organize meetings, somehow keep development going and provide infrastructure and distributions. They are organized around languages. According to Diamond new languages are still discovered but many go extinct too. So the potential for language related user groups is not really growing.

Some of the problems that we face in this world have become too large to be dealt with by individuals and nations. In spite of what anti-globalists want we cannot deal with our energy hunger, environmental issues, lack of natural resources, upcoming technologies without global cooperation. We currently see a regression in cooperation by nationalistic movements, protectionism and the usual going back to presumed better times, but that won't work.

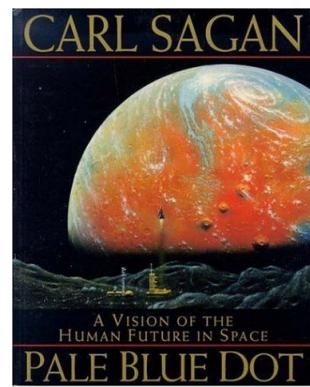
Local user groups are important but the number of members is not growing. There is some cooperation between groups but eventually we might need to combine the groups into one which might succeed unless one wants to come first. Of course we will get the same sentiments and arguments as in regular politics but on the other hand, we already have the advantage of TeX systems being multi-lingual and users sharing interest in the diversity of usage and users. The biggest challenge is to pass on what we have achieved. We're just a momentary highlight and let's not try to embrace some "TeX first" madness.



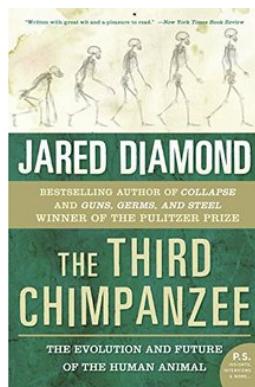
art



history



astronomy



history

Figure 1.4

1.11 Sexes

Most species have two sexes but it is actually a continuum controlled by hormones and genetic expression: we just have to accept it. Although the situation has improved there are plenty of places where some gender relationships are considered bad even to the extent that one's life can be in danger. Actually having strong ideas about these issues is typically human. But in the end one has to accept the continuum.

In a similar way we just have to accept that \TeX usage, application of \TeX engines, etc. is a continuum and not a batch versus wysiwyg battle any more. It's disturbing to read strong recommendations not to use this or that. Of the many macro packages that showed up only a few were able to survive. How do users of outlines look at bitmaps, how do dvi lovers look at pdf. But, as typesetting relates to esthetics, strong opinions come with the game.

Sapolsky reports about a group of baboons where due to the fact that they get the first choice of food the alpha males of pack got poisoned, so that the remaining suppressed males who treated the females well became dominant. In fact they can then make sure that no new alpha male from outside joins the pack without behaving like they do. A sort of social selection. In a similar fashion, until now the gatherings of \TeX ies managed to keep its social properties and has not been dominated by for instance commerce.

In the animal world often sexes relate to appearance. The word sexy made it to other domains as well. Is \TeX sexy? For some it is. We often don't see the real colors of birds. What looks gray to us looks vivid to a bird which sees in a different spectrum. The same is true for \TeX . Some users see a command line (shell) and think: this is great! Others just see characters and keystrokes and are more attracted to an interactive program. When I see a graphic made by MetaPost, I always note how exact it is. Others don't care if their interactive effort doesn't connect the dots well. Some people (also present here) think that we should make \TeX attractive but keep in mind that like and dislike are not fixed human properties. Some mindsets might as well be the result from our makeup, others can be driven by culture.

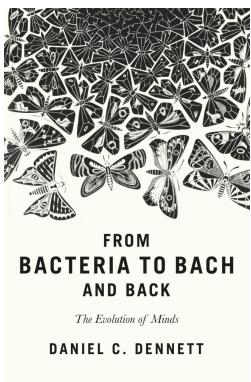
1.12 Religion

One of Sapolsky's lectures is about religion and it comes in the sequence of mental variations including depression and schizophrenia, because all these relate to mental states, emotions, thresholds and such (all things human). That makes it a tricky topic which is why it has not been taped. As I was raised in a moderate Protestant tradition I can imagine that it's an uncomfortable topic instead. But there are actually a few years older videos around and they are interesting to watch and not as threatening as some might expect. Here I just stick to some common characteristics.

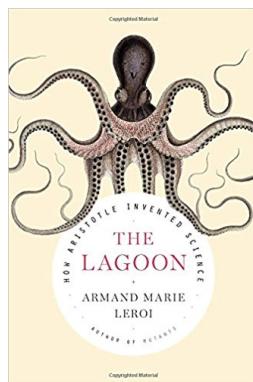
If you separate the functions that religions play into for instance explanation of the yet unknown, social interactions, control of power and regulation of morals, then it's clear why at \TeX user group meetings the religious aspect of \TeX has been discussed in talks. Those who see programs as infallible and always right and don't understand the inner working can see it as an almighty entity. In the Netherlands church-going diminishes but it looks like alternative meetings are replacing it (and I'm not talking of football matches). So what are our \TeX meetings? What do we believe in? The reason that I bring up this aspect is that in the \TeX community we can find aspects of the more extremist aspects of religions: if you don't use the macro package that I use, you're wrong. If you don't use the same operating system as I do, you're evil. You will be punished if you use the wrong editor for \TeX ? Why don't you use this library (which, by the way, just replaced that other one)? We create angels and daemons. Even for quite convinced atheists (it's not hard to run into them on youtube) a religion only survives when it has benefits, something that puzzles them. So when we're religious about \TeX and friends we have to make sure that it's at least beneficial. Also, maybe we fall in Dennett's category of "believers who want to believe": it helps us to do our job if we just believe that we have the perfect tool. Religion has inspired visual and aural art and keeps doing that. (Don Knuth's current musical composition project is a good example of this.)

Scientists can be religious, in flexible ways too, which is demonstrated by Don Knuth. In fact, I'm pretty sure \TeX would not be in the position it is in now if it weren't for his knowledgeable, inspirational, humorous, humble, and always positive presence. And for sure he's not at all religious about the open source software that he sent viral.

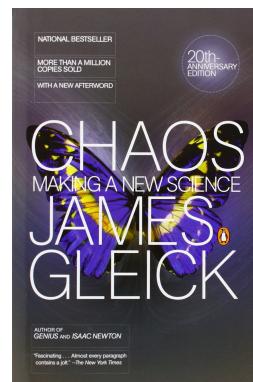
I'm halfway through reading "The Good Book of Human Nature" (An Evolutionary Reading of the Bible) a book about the evolution of the bible and monotheism which is quite interesting. It discusses for instance how transitions from a hunter to a farmer society demanded a change of rules and introduced stories that made sense in that changing paradigm. Staying in one place means that possessions became more important and therefore inheritance. Often when religion is discussed by behavioral biologists, historians and anthropologists they stress this cultural narrative aspect. Also mentioned is that such societies were willing to support (in food and shelter) the ones that didn't normally fit in but added to the spiritual character of religions. The social and welcoming aspect is definitely present in for instance Bacho \TeX conferences although a bystander



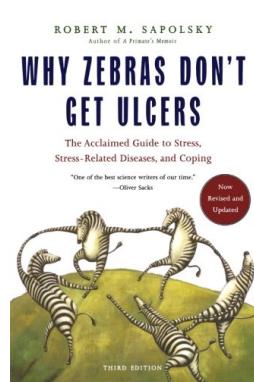
philosophy



science history



science



behavioral
biology

Figure 1.5

can wonder what these folks are doing in the middle of the night around a campfire, singing, drinking, frying sausages, spitting fire, and discussing the meaning of life.

Those who wrap up the state of religious affairs, do predictions and advocate the message, are sometimes called evangelists. I remember a \TeX conference in the usa where the gospel of xml was preached (by someone from outside the \TeX community). We were all invited to believe it. I was sitting in the back of the crowded (!) room and that speaker was not at all interested in who spoke before and after. Well, I do my share of xml processing with ConTeXt , but believe me: much of the xml that we see is not according to any gospel. It's probably blessed the same way as those state officials get blessed when they ask and pray for it in public.

It can get worse at \TeX conferences. Some present here at Bacho \TeX might remember the pdf evangelists that we had show up at \TeX conferences. You see this qualification occasionally and I have become quite allergic to qualifications like architect, innovator, visionary, inspirator and evangelist, even worse when they look young but qualify as senior. I have no problem with religion at all but let's stay away from becoming one. And yes, typography also falls into that trap, so we have to be doubly careful.

1.13 Chaotic solutions

The lectures on “chaos and reductionism” and “emergence and complexity” were the highlights in Sapolsky’s lectures. I’m not a good narrator so I will not summarize them but it sort of boils down to the fact that certain classes of problems cannot be split up in smaller tasks that we understand well, after which we can reassemble the solutions to deal with the complex task. Emerging systems can however cook up working solutions from random events. Examples are colonies of ants and bees.

The \TeX community is like a colony: we cook up solutions, often by trial and error. We dream of the perfect solutions but deep down know that esthetics cannot be programmed in detail. This is a good thing because it doesn’t render us obsolete. At last year’s Bacho \TeX , my nephew Teun and I challenged the anthill outside the canteen to

typeset the \TeX logo with sticks but it didn't persist. So we don't need to worry about competition from that end. How do you program a hive mind anyway?

When chaos theory evolved in the second half of the previous century not every scientist felt happy about it. Instead of converging to more perfect predictions and control in some fields a persistent uncertainty became reality.

After about a decade of using \TeX and writing macros to solve recurring situations I came to the conclusion that striving for a perfect \TeX (the engine) that can do everything and anything makes no sense. Don Knuth not only stopped adding code when he could do what he needed for his books, he also stuck to what to me seems reasonable endpoints. Every hard-coded solution beyond that is just that: a hard-coded solution that is not able to deal with the exceptions that make up most of the more complex documents. Of course we can theorize and discuss at length the perfect never-reachable solutions but sometimes it makes more sense to admit that an able user of a desktop publishing system can do that job in minutes, just by looking at the result and moving around an image or piece of text a bit.

There are some hard-coded solutions and presets in the programs but with LuaTeX and mplib we try to open those up. And that's about it. Thinking that for instance adding features like protrusion or expansion (or whatever else) always lead to better results is just a dream. Just as a butterfly flapping its wings on one side of the world can have an effect on the other side, so can adding a single syllable to your source completely confuse an otherwise clever column or page break algorithm. So, we settle for not adding more to the engine, and provide just a flexible framework.

A curious observation is that when Edward Lorenz ran into chaotic models it was partially due to a restart of a simulation midway, using printed floating point numbers that then in the computer were represented with a different accuracy than printed. Aware of floating point numbers being represented differently across architectures, Don Knuth made sure that \TeX was insensitive to this so that its outcome was predictable, if you knew how it worked internally. Maybe LuaTeX introduces a bit of chaos because the Lua we use has only floats. In fact, a few months ago we did uncover a bug in the backend where the same phenomena gave a chaotic crash.

In chaos theory there is the concept of an attractor. When visualized this can be the area (seemingly random) covered by a trajectory. Or it can be a single point where for instance a pendulum comes to rest. So what is our attractor? We have a few actually. First there is the engine, the stable core of primitives always present. You often see programs grow more complex every update and for sure that happened with e-\TeX , pdfTeX , XeTeX and LuaTeX . However there is always the core that is supposed to be stable. After some time the new kid arrives at a stable state not much different from the parent. The same is true for MetaPost. Fonts are somewhat different because the technology changes but in the end the shapes and their interactions become stable as well. Yet another example is \TeX Live: during a year it might diverge from its route but eventually it settles down and enters the area where we expect it to end up. The \TeX world is at times chaotic, but stable in the long run.

So, how about the existence, the reason for it still being around? One can speculate about its future trajectory but one thing is sure: as long as we break a text into paragraphs and pages \TeX is hard to beat. But what if we don't need that any more? What if the concept of a page is no longer relevant? What if justified texts no longer matter (often designers don't care anyway)? What if students are no longer challenged to come up with a nice looking thesis? Do these collaborative tools with remote \TeX processing really bring new long term users or is \TeX then just one of the come-and-go tools?

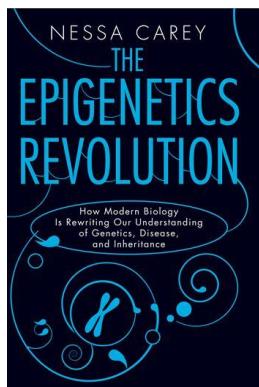
1.14 Looking ahead

In an interview ("World of ideas") Asimov explains that science fiction evolved rapidly when people lived long enough to see that there was a future (even for their offspring) that is different from today. It is (at least for me) mind boggling to think of an evolution of hundreds of thousands of years to achieve something like language. Waiting for the physical being to arrive at a spot where you can make sounds, where the brain is suitable for linguistic patterns, etc. A few hundred years ago speed of any developments (and science) stepped up.

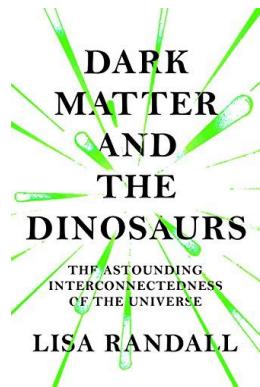
\TeX is getting near 40 years old. Now, for software that **is old!** In that period we have seen computers evolve: thousands of times faster processing, even more increase in memory and storage. If we read about spaceships that travel at a reasonable fraction of the speed of light, and think that will not happen soon, just think back to the terminals that were sitting in computer labs when \TeX was developed: 300 baud was normal. I actually spent quite some time on optimizing time-critical components of ConTeXt but on this timescale that is really a waste of time. But even temporary bottlenecks can be annoying (and costly) enough to trigger such an effort. (Okay, I admit that it can be a challenge, a kind of game, too.)

Neil Tyson, in the video "Storytelling of science" says that when science made it possible to make photos it also made possible a transition in painting to impressionism. Other technology could make the exact snapshot so there was new room for inner feelings and impressions. When the Internet showed up we went through a similar transition, but \TeX actually dates from before the Internet. Did we also have a shift in typesetting? To some extent yes, browsers and real time rendering is different from rendering pages on paper. In what space and time are \TeX ies rooted?

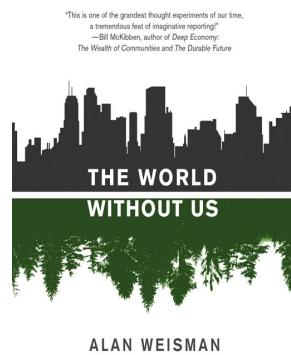
We get older than previous generations. Quoting Sapolsky "... we are now living well enough and long enough to slowly fall apart." The opposite is happening with our tools, especially software: its useful lifetime becomes shorter and changes faster each year. Just look at the version numbers of operating systems. Don Knuth expected \TeX to last for a long time and compared to other software its core concept and implementation is doing surprisingly well. We use a tool that suits our lifespan! Let's not stress ourselves out too much with complex themes. (It helps to read "Why zebras don't get ulcers".)



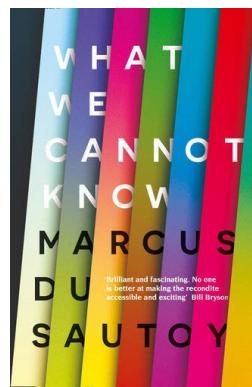
genetics



physics



history



science

Figure 1.6

1.15 Memes

If you repeat a message often enough, even if it's something not true, it can become a meme that gets itself transferred across generations. Conferences like this is where they can evolve. We tell ourselves and the audience how good \TeX is and because we spend so many hours, days, weeks, months using it, it actually must be good, or otherwise we would not come here and talk about it. We're not so stupid as to spend time on something not good, are we? We're always surprised when we run into a (potential) customer who seems to know \TeX . It rings a bell, and it being around must mean something. Somehow the \TeX meme has anchored itself when someone attended university. Even if experiences might have been bad or usage was minimal. The meme that \TeX is the best in math typesetting is a strong survivor.

There's a certain kind of person who tries to get away with their own deeds and decisions by pointing to "fake news" and accusations of "mainstream media" cheating on them. But to what extent are our stories true about how easy \TeX macro packages are to use and how good their result? We have to make sure we spread the right memes. And the user groups are the guardians.

Maybe macro packages are like memes too. In the beginning there was a bunch but only some survived. It's about adaptation and evolution. Maybe competition was too fierce in the beginning. Like ecosystems, organisms and cellular processes in biology we can see the \TeX ecosystem, users and usage, as a chaotic system. Solutions pop up, succeed, survive, lead to new ones. Some look similar and slightly different input can give hugely different outcomes. You cannot really look too far ahead and you cannot deduce the past from the present. Whenever something kicks it off its stable course, like the arrival of color, graphics, font technologies, pdf, xml, ebooks, the \TeX ecosystem has to adapt and find its stable state again. The core technology has proven to be quite fit for the kind of adaptation needed. But still, do it wrong and you get amplified out of existence, don't do anything and the external factors also make you extinct. There is no denial that (in the computer domain) \TeX is surprisingly stable and adaptive. It's also hard not to see how conservatism can lead to extinction.

1.16 Inspiration

I just took some ideas from different fields. I could have mentioned quantum biology, which tries to explain some unexplainable phenomena in living creatures. For instance how do birds navigate without visible and measurable clues. How do people arrive at \TeX while we don't really advertise? Or I could mention epigenetics and explorations in junk dna. It's not the bit of the genome that we thought that matters, but also the expression of the genes driven by other factors. Offspring not only gets genetic material passed but it can get presets. How can the \TeX community pass on Knuth's legacy? Do we need to hide the message in subtle ways? Or how about the quest for dark matter? Does it really exist or do we want (need) it to exist? Does \TeX really have that many users, or do we cheat by adding the users that are enforced during college but don't like it at all? There's enough inspiration for topics at \TeX conferences, we just have to look around us.

1.17 Stability

I didn't go into technical aspects of \TeX yet. I must admit that after decades of writing macros I've reached a point where I can safely say that there will never be perfect automated solutions for really complex documents. When books about neural networks show up I wondered if it could be applied (but I couldn't). When I ran into genetic algorithms I tried to understand its possible impact (but I never did). So I stuck to writing solutions for problems using visualization: the trial and error way. Of course, speaking of ConTeXt, I will adapt what is needed, and others can do that as well. Is there a new font technology? Fine, let's support it as it's no big deal, just a boring programming task. Does a user want a new mechanism? No problem, as solving a reduced subset of problems can be fun. But to think of \TeX in a reductionist way, i.e. solving the small puzzles, and to expect the whole to work in tandem to solve a complex task is not trivial and maybe even impossible. It's a good thing actually, as it keeps us on edge. Also, ConTeXt was designed to help you with your own solutions: be creative.

I mentioned my nephew Bram. He has seen part of this crowd a few times, just like his brother and sister do now. He's into artificial intelligence now. In a few years I'll ask him how he sees the current state of \TeX affairs. I might learn a few tricks in the process.

In "The world without us" Weisman explores how fast the world would be void of traces of humankind. A mere 10.000 years can be more than enough. Looking back, that's about the time hunters became farmers. So here's a challenge: say that we want an ant culture that evolves to the level of having archaeologists to know that we were here at BachoTeX . . . what would we leave behind?

Sapolsky ends his series by stressing that we should accept and embrace individual differences. The person sitting next to you can have the same makeup but be just a bit more sensitive to depression or be the few percent with genes controlling schizophrenic behaviour. He stresses that knowing how things work or where things go wrong doesn't

mean that we should fix everything. So look at this room full of $\text{\TeX}ies$: we don't need to be all the same, use all the same, we don't need some dominance, we just need to accept and especially we need to understand that we can never fully understand (and solve) everything forever.

Predictions, one of the themes, can be hard. It's not true that science has the answer to everything. There will always be room for speculation and maybe we will always need metaphysics too. I just started to read "What we cannot know" by Sautoy. For sure those present here can not predict how \TeX will go on and/or be remembered.

1.18 Children of \TeX

I mentioned "Children of time". The author lets you see their spidery world through spider eyes and physiology. They have different possibilities (eyesight, smell) than we do and also different mental capabilities. They evolve rapidly and have to cope conceptually with signals from a human surveillance satellite up in the sky. Eventually they need to deal with a bunch of (of course) quarrelling humans who want their place on the planet. We humans have some pre-occupation with spiders and other creatures. In a competitive world it is sometimes better to be suspicious (and avoid and flee) than to take a risk of being eaten. A frequently used example is that a rustle in a bush can be the wind or a lion, so best is to run.

We are not that well adapted to our current environment. We evolved at a very slow pace so there was no need to look ahead more than a year. And so we still don't look too far ahead (and choose politicians accordingly). We can also not deal that well with statistics (Dawkins's "Climbing Mount Probability" is a good read) so we make false assumptions, or just forget.

Does our typeset text really look that good on the long run, or do we cheat with statistics? It's not too hard to find a bad example of something not made by \TeX and extrapolate that to the whole body of typeset documents. Just like we can take a nice example of something done by \TeX and assume that what we do ourselves is equally okay. I still remember the tests we did with $\text{pdf}\text{\TeX}$ and hz . When Han The Thanh and I discussed that with Hermann Zapf he was not surprised at all that no one saw a difference between the samples and instead was focusing on aspects that $\text{\TeX}ies$ are told to look at, like two hyphens in a row.

A tool like \TeX has a learning curve. If you don't like that just don't use it. If you think that someone doesn't like that, don't enforce this tool on that someone. And don't use (or lie with) statistics. Much better arguments are that it's a long-lived stable tool with a large user base and support. That it's not a waste of time. Watching a designer like Hermann Zapf draw shapes is more fun than watching click and point in heavily automated tools. It's probably also less fun to watch a $\text{\TeX}ie$ converge towards a solution.

Spiders are resilient. Ants maybe even more. Ants will survive a nuclear blast (mutations might even bring them benefits), they can handle the impact of a meteorite, a

change in climate won't harm them much. Their biggest enemy is probably us, when we try to wipe them out with poison. But, as long as they keep a low profile they're okay. \TeX doesn't fit into the economic model as there is no turnaround involved, no paid development, it is often not seen at all, it's just a hit in a search engine and even then you might miss it (if only because no one pays for it being shown at the top).

We can learn from that. Keeping a low profile doesn't trigger the competition to wipe you out. Many (open source) software projects fade away: some big company buys out the developer and stalls the project or wraps what they bought in their own stuff, other projects go professional and enterprise and alienate the original users. Yet others abort because the authors lose interest. Just like the ideals of socialism don't automatically mean that every attempt to implement it is a success, so not all open source and free software is good (natured) by principle either. The fact that communism failed doesn't mean that capitalism is better and a long term winner. The same applies to programs, whether successful or not.

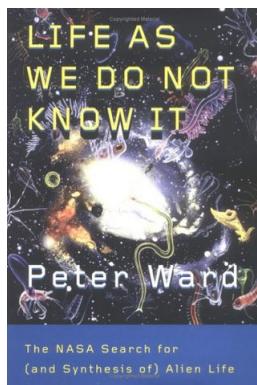
Maybe we should be like the sheep. Dennett uses these animals as a clever species. They found a way to survive by letting themselves (unconsciously) be domesticated. The shepherd guarantees food, shelter and protection. He makes sure they don't get ill. Speaking biologically: they definitely made sure that many copies of their genes survived. Cows did the same and surprisingly many of them are related due to the fact that they share the same father (something now trying to be reverted). All \TeX spin-offs relate to the same parent, and those that survived are those that were herded by user groups. We see bits and pieces of \TeX end up in other applications. Hyphenation is one of them. Maybe we should settle for that small victory in a future hall of fame.

When I sit on my balcony and look at the fruit trees in my garden, some simple math can be applied. Say that one of the apple trees has 100 apples per year and say that this tree survives for 25 years (it's one of those small manipulated trees). That makes 2.500 apples. Without human intervention only a few of these apples make it into new trees, otherwise the whole world would be dominated by apple trees. Of course that tree now only survives because we permit it to survive, and for that it has to be humble (something that is very hard for modern Apples). Anyway, the apple tree doesn't look too unhappy.

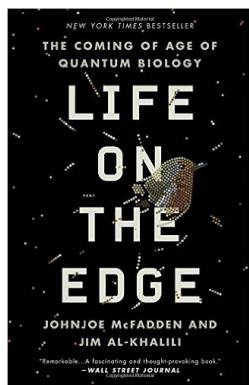
A similar calculation can be done for birds that nest in the trees and under my roof. Given that the number of birds stays the same, most of energy spent on raising offspring is wasted. Nevertheless they seem to enjoy life. Maybe we should be content if we get one enthusiastic new user when we demonstrate \TeX to thousands of potential users.

Maybe, coming back to the themes of the conference, we should not come up with these kinds of themes. We seem to be quite happy here. Talking about the things that we like, meeting people. We just have to make sure that we survive. Why not stay low under the radar? That way nothing will see us as a danger. Let's be like the ants and spiders, the invisible hive mind that carries our message, whatever that is.

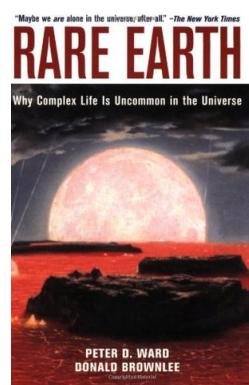
When Dennett discusses language he mentions (coined) words that survive in language. He also mentions that children pick up language no matter what. Their minds



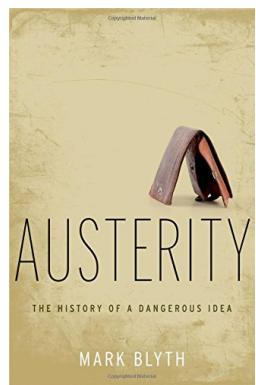
astrobiology



quantumbiology



astrophysics



economics

Figure 1.7

are made for it. Other animals don't do that: they listen but don't start talking back. Maybe \TeX is just made for certain minds. Some like it and pick it up, while for others it's just noise. There's nothing wrong with that. Predilection can be a user property.

1.19 The unexpected

In a discussion with Dawkins the well-spoken astrophysicist Neil deGrasse Tyson brings up the following. We differ only a few percent in dna from a chimp but quite a lot in brain power, so how would it be if an alien that differs a few percent (or more) passes by earth. Just like we don't talk to ants or chimps or whatever expecting an intelligent answer, whatever passes earth won't bother wasting time on us. Our rambling about the quality of typesetting probably sounds alien to many people who just want to read and who happily reflow a text on an ebook device, not bothered by a lack of quality.

We tend to take ourselves as reference. In "Rare Earth" Ward and Brownlee extrapolate the possibility of life elsewhere in the universe. They are not alone in thinking that while on one hand applying statistics to these formulas of possible life on planets there might also be a chance that we're the only intelligent species ever evolved. In a follow up, "Life as we do not know it" paleontologist and astrobiologist Ward (one of my favourite authors) discusses the possibility of life not based on carbon, which is not natural for a carbon based species. Carl Sagan once pointed out that an alien species looking down to earth can easily conclude that cars are the dominant species on earth and that the thingies crawling in and out them are some kind of parasites. So, when we look at the things that somehow end up on paper (as words, sentences, ornaments, etc.), what is dominant there? And is what we consider dominant really that dominant in the long run? You can look at a nice page as a whole and don't see the details of the content. Maybe beauty hides nonsense.

When \TeX ies look around they look to similar technologies. Commands in shells and solutions done by scripting and programming. This make sense in the perspective of survival. However, if you want to ponder alternatives, maybe not for usage but just for fun, a completely different perspective might be needed. You must be willing to accept that communicating with a user of a wysiwyg program might be impossible. If mutual

puzzlement is a fact, then they can either be too smart and you can be too dumb or the reverse. Or both approaches can be just too alien, based on different technologies and assumptions. Just try to explain \TeX to a kid 40 years younger or to an 80 year old grandparent for that matter. Today you can be very clever in one area and very stupid in another.

In another debate, Neil deGrasse Tyson asks Dawkins the question why in science fiction movies the aliens look so human and when they don't, why they look so strange, for instance like cumbersome sluggish snails. The response to that is one of puzzlement: the opponent has no reference of such movies. In discussions old \TeX ies like to suggest that we should convert young users. They often don't understand that kids live in a different universe.

How often does that happen to us? In a world of many billions \TeX has its place and can happily coexist with other typesetting technologies. Users of other technologies can be unaware of us and even create wrong images. In fact, this also happens in the community itself: (false) assumptions turned into conclusions. Solutions that look alien, weird and wrong to users of the same community. Maybe something that I present as hip and modern and high- \TeX and promising might be the opposite: backward, old-fashioned and of no use to others. Or maybe it is, but the audience is in a different mindset. Does it matter? Let's just celebrate that diversity. (So maybe, instead of discussing the conference theme, I should have talked about how I abuse \LaTeX in controlling lights in my home as part of some IoT experiments.)

1.20 What drives us

I'm no fan of economics and big money talk makes me suspicious. I cannot imagine working in a large company where money is the drive. It also means that I have not much imagination in that area. We get those calls at the office from far away countries who are hired to convince us by phone of investments. Unfortunately mentioning that you're not at all interested in investments or that multiplying money is irrelevant to you does not silence the line. You have to actively kill such calls. This is also why I probably don't understand today's publishing world where money also dominates. Recently I ran into talks by Mark Blyth about the crisis (what crisis?) and I wish I could argue like he does when it comes to typesetting and workflows. He discusses quite well that most politicians have no clue what the crisis is about.

I think that the same applies to the management of publishers: many have no clue what typesetting is about. So they just throw lots of money into the wrong activities, just like the central banks seem to do. It doesn't matter if we \TeX ies demonstrate cheap and efficient solutions.

Of course there are exceptions. We're lucky to have some customers that do understand the issues at hand. Those are also the customers where authors may use the tools themselves. Educating publishers, and explaining that authors can do a lot, might be a premise, predilection and prediction in one go! Forget about those who don't get it:

they will lose eventually, unfortunately not before they have reaped and wasted the landscape.

Google, Facebook, Amazon, Microsoft and others invest a lot in artificial intelligence (or, having all that virtual cash, just buy other companies that do). They already have such entities in place to analyze whatever you do. It is predicted that at some point they know more about you than you know yourself. Reading Luke Dormehl's "The Formula" is revealing. So what will that do with our so-called (disputed by some) free will? Can we choose our own tools? What if a potential user is told that all his or her friends use WhateverOffice so they'd better do that too? Will subtle pressure lead them or even us users away from \TeX ? We already see arguments among \TeX ies, like "It doesn't look updated in 3 years, is it still good?" Why update something that is still valid? Will the community be forced to update everything, sort of fake updates. Who sets out the rules? Do I really need to update (or re-run) manuals every five years?

Occasionally I visit the Festo website. This is a (family owned) company that does research at the level that used to be common in large companies decades ago. If I had to choose a job, that would be the place to go to. Just google for "festo bionic learning network" and you understand why. We lack this kind of research in the field we talk about today: research not driven by commerce, short term profit, long term control, but because it is fundamental fun.

Last year Alan Braslau and I spent some time on bib \TeX . Apart from dealing with all the weird aspects of the apa standard, dealing with the inconsistently constructed author fields is a real pain. There have been numerous talks about that aspect here at Bacho \TeX by Jean-Michel Hufflen. We're trying to deal with a more than 30-year-old flawed architecture. Just look back over a curve that backtracks 30 years of exponential development in software and databases and you realize that it's a real waste of time and a lost battle. It's fine to have a text based database, and stable formats are great, but the lack of structure is appalling and hard to explain to young programmers. Compare that to the Festo projects and you realize that there can be more challenging projects. Of course, dealing with the old data can be a challenge, a necessity and eventually even be fun, but don't even think that it can be presented as something hip and modern. We should be willing to admit flaws. No wonder that Jean-Michel decided to switch to talking about music instead. Way more fun.

Our brains are massively parallel bio-machinery. Groups of neurons cooperate and compete for attention. Coming up with solutions that match what comes out of our minds demands a different approach. Here we still think in traditional programming solutions. Will new ideas about presenting information, the follow up on books come from this community? Are we the innovative Festo or are we an old dinosaur that just follows the fashion?

1.21 User experience

Here is a nice one. Harari spends many pages explaining that research shows that when an unpleasant experience has less unpleasantness at the end of the period involved, the

overall experience is valued according to the last experience. Now, this is something we can apply to working with \TeX : often, the more you reach the final state of typesetting the more it feels as all hurdles are in the beginning: initial coding, setting up a layout, figuring things out, etc.

It can only get worse if you have a few left-over typesetting disasters but there adapting the text can help out. Of course seeing it in a cheap bad print can make the whole experience bad again. It happens. There is a catch here: one can find lots of bad-looking documents typeset by \TeX . Maybe there frustration (or indifference) prevails.

I sometimes get to see what kind of documents people make with Con \TeX t and it's nice to see a good looking thesis with diverse topics: science, philosophy, music, etc. Here \TeX is just instrumental, as what it is used for is way more interesting (and often also more complex) than the tool used to get it on paper. We have conferences but they're not about rocket science or particle accelerators. Proceedings of such conferences can still scream \TeX , but it's the content that matters. Here somehow \TeX still sells itself, being silently present in rendering and presentations. It's like a rootkit: not really appreciated and hard to get rid of. Does one discuss the future of rootkits other than in the perspective of extinction? So, even as an invisible rootkit, hidden in the workings of other programs, \TeX 's future is not safe. Sometimes, when you install a Linux system, you automatically get this large \TeX installation, either because of dependencies or because it is seen as a similar toolkit as for instance Open (or is it Libre) Office. If you don't need it, that user might as well start seeing it as a (friendly) virus.

1.22 Conclusion

At some point those who introduced computers in typesetting had no problem throwing printing presses out of the window. So don't pity yourself if at some point in the near future you figure out that professional typesetting is no longer needed. Maybe once we let machines rule the world (even more) we will be left alone and can make beautiful documents (or whatever) just for the joy, not bothering if we use outdated tools. After all, we play modern music on old instruments (and the older rock musicians get, the more they seem to like acoustic).

There are now computer generated compositions that experienced listeners cannot distinguish from old school. We already had copies of paintings that could only be determined forgeries by looking at chemical properties. Both of these (artificial) arts can be admired and bring joy. So, the same applies to fully automated typeset novels (or runtime rendered ebooks). How bad is that really? You don't dig channels with your hand. You don't calculate logarithmic tables manually any longer.

However, one of the benefits of the Internet is watching and listening to great minds. Another is seeing musicians perform, which is way more fun than watching a computer (although googling for "animusic" brings nice visuals). Recently I ran into a wooden musical computer made by "Wintergatan" which reminded me of the "Paige Composer" that we use in a Lua \TeX cartoon. Watching something like that nicely compensates

for a day of rather boring programming. Watching how the marble machine x (mmx) evolves is yet another nice distraction.

Now, the average age of the audience here is pretty high even if we consider that we get older. When I see solutions of ConTeXt users (or experts) posted by (young) users on the mailing list or stack exchange I often have to smile because my answer would have been worse. A programmable system invokes creative solutions. My criterion is always that it has to look nice in code and has some elegance. Many posted solutions fit. Do we really want more automation? It's more fun to admire the art of solutions and I'm amazed how well users use the possibilities (even ones that I already forgot).

One of my favourite artists on my weekly "check youtube" list is Jacob Collier. Right from when I ran into him I realized that a new era in music had begun. Just google for his name and "music theory interview" and you probably understand what I mean. When Dennett comments on the next generation (say up to 25) he wonders how they will evolve as they grow up in a completely different environment of connectivity. I can see that when I watch family members. Already long ago Greg Bear wrote the novel "Darwin's Children". It sets you thinking and when looking around you even wonder if there is a truth in it.

There are folks here at BachoTeX who make music. Now imagine that this is a conference about music and that the theme includes the word "future". Then, imagine watching that video. You see some young musicians, one of them probably one of the musical masterminds of this century, others instrumental to his success, for instance by wrapping up his work. While listening you realize that this next generation knows perfectly well what previous generations did and achieved and how they influenced the current. You see the future there. Just look at how old musicians reflect on such videos. (There are lots of examples of youth evolving into prominent musicians around and I love watching them). There is no need to discuss the future, in fact, we might make a fool of ourselves doing so. Now back to this conference. Do we really want to discuss the future? What we think is the future? Our future? Why not just hope that in the flow of getting words on a medium we play our humble role and hope we're not forgotten but remembered as inspiration.

One more word about predicting the future. When Arthur Clarke's "2001: A Space Odyssey" was turned into a movie in 1968, a lot of effort went into making sure that the not so far ahead future would look right. In 1996 scientists were asked to reflect on these predictions in "Hal's Legacy". It turned out that most predictions were plain wrong. For instance computers got way smaller (and even smaller in the next 20 years) while (self-aware) artificial intelligence had not arrived either. So, let's be careful in what we predict (and wish for).

1.23 No more themes

We're having fun here, that's why we come to BachoTeX (predilection). That should be our focus. Making sure that TeX's future is not so much in the cutting edge but in

providing fun to its users (prediction). So we just have to make sure it stays around (premise). That's how it started out. Just watch at Don Knuth's 3:16 poster: via \TeX and MetaFont he got in contact with designers and I wouldn't be surprised if that sub-project was among the most satisfying parts. So, maybe instead of ambitious themes the only theme that matters is: show what you did and how you did it.

Advertising TEX 2

I can get upset when I hear TEXies boast about the virtues of TEX compared to for instance Microsoft Word. Not that I feel responsible for defending a program that I never use(d) but attacking something for no good reason makes not much sense to me. It is especially annoying when the attack is accompanied by a presentation that looks pretty bad in design and typography. The best advertisements for TEX should of course come from outside the TEX community, by people impressed by its capabilities. How many TEXies can really claim that Word is bad when they never tried to make something in it with a similar learning curve as they had in TEX or the same amount of energy spent in editing and perfecting a word-processor-made document.

In movies where computer technology plays a role one can encounter weird assumptions about what computers and programs can do. Run into a server room, pull one disk out of a raid-5 array and get all information from it. Connect some magic device to a usb port of a phone and copy all data from it in seconds. Run a high speed picture or fingerprint scan on a computer (probably on a remote machine) and show all pictures flying by. Okay, it's not so far from other unrealistic aspects in movies, like talking animals, so maybe it is just a metaphor for complexity and speed. When zapping channels on my television I saw figure 2.1 and as the media box permits replay I could make a picture. I have no clue what the movie was about or what movie it was so a reference is lacking here. Anyway it's interesting that seeing a lot of TEX code flying by can impress someone: the viewer, even if no TEXie will ever see that on the console unless in some error or tracing message and even then it's hard to get that amount. So, the viewer will never realize that what is seen is definitely not what a TEXie wants to see.

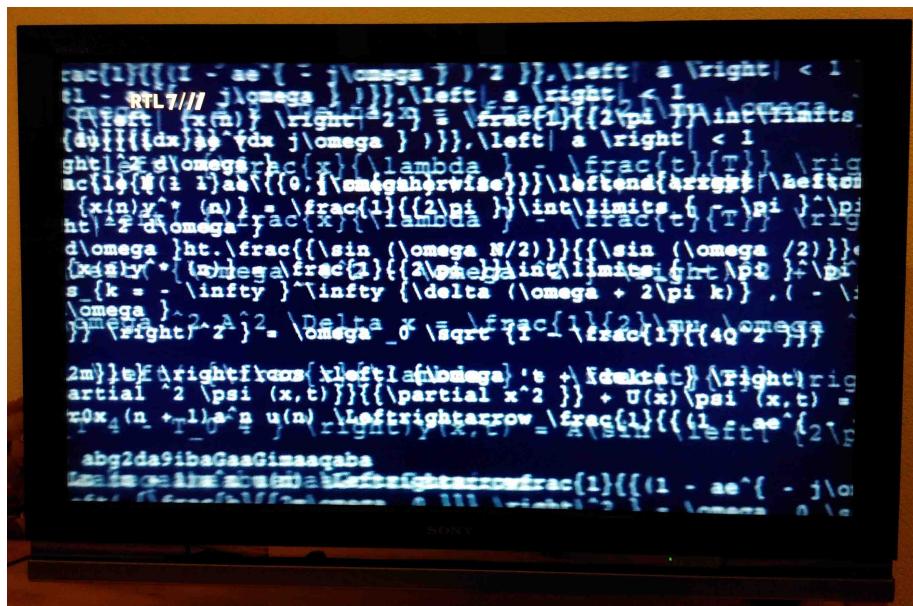


Figure 2.1 TEX in a movie

So, as that kind of free advertisement doesn't promote TeX well, what of an occasional mentioning of TeX in highly-regarded literature? When reading "From bacteria to Bach and back, the evolution of minds" by Daniel Dennett I ran into the following:

"In Microsoft Word, for instance, there are the typographical operations of superscript and subscript, as illustrated by

base^{power}

and

human_{female}

But try to add another superscript to $\text{base}^{\text{power}}$ —it *should* work, but it doesn’t! In mathematics, you can raise powers to powers to powers forever, but you can’t get Microsoft Word to display these (there are other text-editing systems, such as TeX, that can). Now, are we sure that human languages make use of true recursion, or might some or all of them be more like Microsoft Word? Might our interpretation of grammars as recursive be rather an elegant mathematical idealization of the actual “moving parts” of a grammar?”

Now, that book is a wonderfully interesting read and the author often refers to other sources. When one reads some reference (with a quote) then one assumes that what one reads is correct, and I have no reason to doubt Dennett in this. But this remark about TeX has some curious inaccuracies.¹

First of all a textual raise or lower is normally not meant to be recursive. Nesting would have interesting consequences for the interline space so one will avoid it whenever possible. There are fonts that have superscript and subscript glyphs and even Unicode has slots for a bunch of characters. I'm not sure what Word does: take the special glyph or use a scaled down copy?

Then there is the reference to \TeX where we can accept that the “E” is not lowered but just kept as a regular “e”. Actually the mentioning of nested scripts refers to typesetting math and that’s what the superscripts and subscripts are for in \TeX . In math mode however, one will normally raise or lower symbols and numbers, not words: that happens in text mode.

While Word will use the regular text font when scripting in text mode, a \TeX user will either have to use a macro to make sure that the right size (and font) is used, or one can revert to math mode. But how to explain that one has to enter math and then explicitly choose the right font? Think of this:

efficient\high{efficient} or
efficient\\$^{\text{efficient}}\\$ or \par

¹ Of course one can wonder in general that when one encounters such an inaccuracy, how valid other examples and conclusions are. However, consistency in arguments and confirmation by other sources can help to counter this.

```
{\bf efficient} high{efficient} or  
efficient$^{\text{efficient}}$
```

Which gives (in Cambria)

efficient^{efficient} or efficient^{efficient} or

efficient^{efficient} or efficient^{efficient}

Now this,

```
efficient high{efficient} or  
efficient$^{\text{efficient}}$ or \par  
{\bf efficient} high{efficient} or  
efficient$^{\text{efficient}}$
```

will work okay but the math variant is probably quite frightening at a glance for an average Word user (or beginner in \TeX) and I can understand why someone would rather stick to click and point.

efficient^{efficient} or efficient^{efficient} or

efficient^{efficient} or efficient^{efficient}

Oh, and it's tempting to try the following:

```
efficient{\addff{f:superiors}efficient}
```

but that only works with fonts that have such a feature, like Cambria:

efficient^{efficient}

To come back to Dennett's remark: when typesetting math in Word, one just has to switch to the math editing mode and one can have nested scripts! And, when using \TeX one should not use math mode for text scripts. So in the end in both systems one has to know what one is doing, and both systems are equally capable.

The recursion example is needed in order to explain how (following recent ideas from Chomsky) for modern humans some recursive mechanism is needed in our wetware. Now, I won't go into details about that (as I can only mess up an excellent explanation) but if you want to refer to \TeX in some way, then expansion² of (either combined or not) snippets of knowledge might be a more interesting model than recursion, because much of what \TeX is capable of relates to expansion. But I leave that to others to explore.³

Now, comparing \TeX to Word is always kind of tricky: Word is a text editor with typesetting capabilities and \TeX is a typesetting engine with programming capabilities. Recursion is not really that relevant in this perspective. Endless recursion in scripts makes

² Expanding macros actually works well with tail recursion.

³ One quickly starts thinking of how `expandafter`, `noexpand`, `unexpanded`, `protected` and other primitives can be applied to language, understanding and also misunderstanding.

little sense and even \TeX has its limits there: the \TeX math engine only distinguishes three levels (text, script and scriptscript) and sometimes I'd like to have a level more. Deeper nesting is just more of scriptscript unless one explicitly enforces some style. So, it's recursive in the sense that there can be many levels, but it also sort of freezes at level three.



Figure 2.2 Nicer than \TeX

I love \TeX and I like what you can do with it and it keeps surprising me. And although mathematics is part of that, I seldom have to typeset math myself. So, I can't help that figure 2.2 impresses me more. It even has the so-familiar-to- \TeX ies dollar symbols in it: the poem "Poetry versus Orchestra" written by Hollie McNish, music composed by Jules Buckley and artwork by Martin Pyper (I have the dvd but you can also find it on YouTube). It reminds me of Don Knuth's talk at a tug meeting. In *Tugboat* 31:2 (2010) you can read Don's announcement of his new typesetting engine $i\text{\TeX}$: "Output can be automatically formatted for lasercutters, embroidery machines, 3D printers, milling machines, and other cnc devices . . .". Now that is something that Word can't do!

Why use \TeX ? 3

Hans Hagen
Hasselt NL
July 2021 (public 2023)

3.1 Introduction

Let's assume that you know what \TeX is: a program that interprets a language with the same name that makes it possible to convert (tagged) input into for instance pdf. For many of its users it is a black box: you key in some text, hit a button and get some typeset result in return. After a while you start tweaking this black box, meet other users (on the web), become more fluent and stick to it forever.

But now let's assume that you don't know \TeX and are in search of a system that helps you create beautiful documents in an efficient way. When your documents have a complex structure you are probably willing to spend some time on figuring out what the best tool is. Even if a search lets you end up with something called \TeX , a three letter word with a dropped E, you still don't know what it is. It helps to search for $\backslash\text{\TeX}$ which is pronounced as **tech**. Advertisement for \TeX is often pretty weak. It's rather easy to point to the numerous documents that can be found on the web. But what exactly does \TeX do and what are its benefits? In order to answer this we need to know who you are: an author, editor, an organization that deals with documents or needs to generate readable output, like publishers do.

3.2 Authors

We start with authors. Students of sciences that use mathematics don't have much of a choice. But most of these documents hardly communicate the message that "Everyone should use \TeX ." or that "All documents produced by \TeX look great." but they do advocate that for rendering math it is a pretty good system. The source code of these documents often look rather messy and unattractive and for a non-math user it can be intimidating. Choosing some lightweight click-and-ping alternative looks attractive.

Making \TeX popular is not going to happen by convincing those who have to write an occasional letter or report. They should just use whatever suits them. On the other hand if you love consistency, long term support, need math, are dealing with a rare language or script, like to reuse content, prefer different styling from one source, use one source for multiple documents, or maybe love open source tools, then you are a candidate. Of course there is a learning curve but normally you can master \TeX rather fast and once you get the hang of it there's often no way back. But you always need to invest a bit beforehand.

So what authors are candidates for \TeX ? It could be that \TeX is the only tool that does the job. If so, you probably learned that from someone who saw you struggle or had the same experience and wrote or talked about it somewhere. In that case using \TeX for creating just one document (like a thesis) makes sense. Otherwise, you should really wonder if you want to invest time in a tool that you probably have to ditch later on as most organizations stick to standard (commercial) word processing tools.

Talking to customers we are often surprised that people have heard about \TeX , or even used it for a few documents in college. Some universities just prescribe the use of \TeX for reporting, so not much of a choice there. Memories are normally rather positive in the sense that they know that it can do the job and that it's flexible.

User group journals, presentations at \TeX meetings, journals, books and manuals that come with \TeX macro packages can all be used to determine if this tool suits an author. Actually, I started using \TeX because the original $\text{\TeX}book$ had some magic, and reading it was just that: reading it, as I had no running implementation. A few years later, when I had to write (evolving) reports, I picked up again. But I'm not a typical user.

3.3 Programmers

When you are a programmer who has to generate reports, for instance in pdf, or write manuals, then \TeX can really be beneficial. Of course \TeX is not always an obvious choice, but if you're a bit able to use it it's hard to beat in quality, flexibility and efficiency. I'm often surprised that companies are willing to pay a fortune for functionality that basically comes for free. Programmers are accustomed to running commands and working in a code editor with syntax highlighting so that helps too. They also recognize when something can be done more efficiently.

When you need to go from some kind of input (document source, database, generated) to some rendered output there currently are a few endpoints: a (dynamic) html page, a pdf document, something useable in a word processor, or a representation using the desktop user interface. It's the second category where \TeX is hard to beat but even using \TeX and MetaPost for creating a chart can make sense.

There are of course special cases where \TeX fits in nicely. Say that you have to combine pdf documents. There are numerous tools to do that and \TeX is one. The advantage of \TeX over other tools is that it's trivial to add additional text, number pages, provide headers and footers. And it will work forever. Why? Because \TeX has been around for decades and will be around for decades to come. It's an independent component. The problem with choosing for \TeX is that the starting point is important. The question is not "What tool should I use?" but "What problem do I need to solve?". An open discussion about the objectives and possibilities is needed, not some checklist based on assumptions. If you don't know \TeX and have never worked with a programmable typesetting environment, you probably don't see the possibilities. In fact, you might even choose for \TeX for the wrong reasons.

The problem with this category of users is that they seldom have the freedom to choose their tools. There are not that many jobs where the management is able to recognize the clever programmer who can determine that \TeX is suitable for a lot of jobs and can save money and time. Even the long term availability and support is not an argument since not only most tools (or even apis) changes every few years but also organizations themselves change ownership, objectives, and personnel on a whim. The concept of ‘long term’ is hard to grasp for most people (just look at politics) and it’s only in retrospect that one can say ‘We used that toolkit for over a decade.’

3.4 Organizations

Authors (often) have the advantage that they can choose themselves: they can use what they like. In practice any decent programmer is able to find the suitable tools but convincing the management to use one of them can be a challenge. Here we’re also talking of ‘comfort zones’: you have to like a tool(chain). Organizations normally don’t look for \TeX . Special departments are responsible for choosing and negotiating whatever is used in a company. Unfortunately companies don’t always start from the open question “We have this problem, we want to go there, what should we do?” and then discuss options with for instance those who know \TeX . Instead requirements are formulated and matches are found. The question then is “Are these requirements cut in stone?” and if not (read: we just omit some requirements when most alternatives don’t meet them), were other requirements forgotten? Therefore organizations can end up with the wrong choice (using \TeX in a situation where it makes no sense) or don’t see opportunities (not using \TeX while it makes most sense). It doesn’t help that a hybrid solution (use a mix of \TeX and other tools) is often not an option. Where an author can just stop using a tool after a few days of disappointment, and where a programmer can play around a bit before making a choice, an organization probably best can start small with a proof of concept.

Let’s take a use case. A publisher wants to automatically convert xml files into pdf. One product can come from multiple sources (we have cases where thousands of small xml files combine into one final product). Say that we have three different layouts: a theory book, a teachers manual and an answer book. In addition special proofing documents have to be rendered. The products might be produced on demand with different topics in any combination. There is at least one image and table per page, but there can be more. There are color and backgrounds used, tables of contents generated, there is extensive cross referencing and an index. Of course there is math.

Now let’s assume an initial setup costs 20K Euro and, what happens often when the real products show up, a revision after one year takes the same amount. We also assume 10K for the following eight years for support. So, we end up with 120K over 10 years. If one goes cheap we can consider half of that, or we can be pessimistic and double the amount.

The first year 10K pages are produced, the second year 20K and after that 30K per year. So, we’re talking of 270K pages. If we include customer specific documents and proofing we might as well end up with a multiple of that.

So, we have 120K Euro divided by 270K pages or about half an Euro per page. But likely we have more pages so it costs less. If we double the costs then we can assume that some major changes took place which means more pages. In fact we had projects where the layout changed, all documents were regenerated and the costs were included in the revision, so far from double. We also see many more pages being generated so in practice the price per page drops below half an Euro. The more we process the cheaper it gets and one server can produce a lot of pages!

Now, the interesting bit of such a calculation is that the costs only concern the hours spent on a solution. A \TeX based system comes for free and there are no license costs. Whatever alternative is taken, even if it is as flexible, it will involve additional costs. From the perspective of costs it's very hard to beat \TeX . Add to that the possibility for custom extensions, long term usage and the fact that one can adapt the system. The main question of course is: does it do the job. The only way to find out is to either experiment (which is free), consult an expert (not free, but then needed anyway for any solution) or ask an expert to make a proof of concept (also not free but relatively cheap and definitely cheaper than a failure). In fact, before making decisions about what solution is best it might be a good idea to check with an expert anyway, because more or less than one thinks might be possible. Also, take into account that the \TeX ecosystem is often one of the first to support new technologies, and normally does that within its existing interface. And there is plenty of free support and knowledge available once you know how to find it. Instead of wasting time and money on advertisement and fancy websites, effort goes into support and development. Even if you doubt that the current provider is around in the decade to come, you can be sure that there will be others, simply because \TeX attracts people. Okay, it doesn't help that large companies like to out source to far-far-away and expect support around the corner, so in the end they might kill their support chain.

When talking of \TeX used in organizations we tend to think of publishers. But this is only a small subset of organizations where information gets transformed into something presentable. For small organizations the choice for \TeX can be easy: costs, long term stability, knowing some experts are driving forces. For large organizations these factors seem (at least to us) hardly relevant. We've (had) projects where actually the choice for using a \TeX based solution was (in retrospect) a negative one: there was no other tool than this relatively unknown thing called \TeX . Or, because the normal tools could not be used, one ended up with a solution where (behind the scenes) \TeX is used, without the organization knowing it. Or, it happened that the problem at hand was mostly one that demands in-depth knowledge of manipulating content, cleaning up messy data, combining resources (images or pdf documents), all things that happen to be available in the perspective of \TeX . If you can solve a hard to solve problem for them then an organization doesn't care what tool you use. What does matter is that the solution runs forever, that costs are controllable and above all, that it "Just works." And if you can make it work fast, that helps too. We can safely claim that when \TeX is evaluated as being a good option, that in the end it always works out quite well.

Among arguments that (large) organizations like to use against a choice for \TeX (or something comparable) are the size of the company that they buy their solution from,

the expected availability for support, and the wide-spread usage of the tool at hand. One can wonder if it also matters that many vendors change ownership, change products every few years, change license conditions when they like, charge a lot for support or just abort a tool chain. Unfortunately when that happens those responsible for choosing such a system can have moved on to another job, so this is seldom part of an evaluation. For the supplier the other side of the table is just as much of a gamble. In that respect, an organization that wants to use an open source (and/or free) solution should realize that getting a return on investment on such a development is pretty hard to achieve. So, who really takes the risk for writing open source?

For us, the reason to develop ConTeXt and make it open is that it fits in our philosophy and we like the community. It is actually not really giving us an advantage commercially: it costs way more to develop, support and keep up-to-date than it will ever return. We can come up with better, faster and easier solutions and in the end we pay the price because it takes less time to cook up styles. So there is some back slash involved because commercially a difficult solution leads to more billable hours. Luckily we tend to avoid wasting time so we improve when possible and then it ends up in the distributed code. And, once the solution is there, anyone can use it. Basically also for us it's just a tool, like the operating system, editor and viewer are. So, what keeps development going is mostly the interaction with the community. This also means that a customer can't really demand functionality for free: either do it yourself, wait for it to show up, or pay for it (which seldom happens). Open source is not equivalent with "You get immediately what you want because someone out there writes the code.". There has to be a valid reason and often it's just users and meetings or just some challenge that drives it.

This being said, it is hard to convince a company to use \TeX . It has to come from users in the organization. Or, what we sometimes see with publishers, it comes with an author team or acquired product line where it's the only option. Even then we seldom see transfer to other branches in the organizations. No one seems to wonder "How on earth can that xml to pdf project produce whatever output in large quantities in a short period of time" while other (past) projects failed. It probably relates to the abstraction of the process. Even among \TeX users it can be that you demonstrate something with a click on a button and that many years afterwards someone present at that moment tells you that they just discovered that this or that can be done by hitting a button. I'm not claiming that \TeX is the magic wand for everything but in some areas it's pretty much ahead of the pack. Go to a \TeX user meeting and you will be surprised about the accumulated diverse knowledge present in the room. It's user demand that drives ConTeXt development, not commerce.

3.5 Choosing

So, where can one find information about \TeX and friends? On the web one has to use the right search keys, so adding `tex` helps: `context tex` or `xml tex pdf` and so on. Can one make a fancy hip website, sure, but it being a life-long, already old and mature environment, and given that it comes for free, or is used low-budget, not much effort

and money can be spent on advertising it. A benefit is that no false promises and hypes are made either. If you want to know more, just ask the right folks.

For all kind of topics one can find interesting videos and blogs. One can subscribe to channels on YouTube or join forums. Unfortunately not that many bloggers or vloggers or podcasters come up with original material every time, and often one starts to recognize patterns and will get bored by repetition of wisdom and arguments. The same is true for manuals. Is a ten year old manual really obsolete? Should we just recompile it to fake an update while in fact there has been no need for it? Should we post twenty similar presentations while one can do? (If one already wants to present the same topic twenty times in the first place?) Maybe one should compare \TeX with cars: they became better over time and can last for decades. And no new user manual is needed.

As with blogs and vlogs advertising \TeX carries the danger for triggering political discussions and drawing people into discussions that are not pleasant: \TeX versus some word processor, open versus closed source, free versus paid software, this versus that operating system, editor such or editor so.

To summarize, it's not that trivial to come up with interesting information about \TeX , unless one goes into details that are beyond the average user. And those who are involved are often involved for a long time so it gets more complex over time. User group journals that started with tutorials later on became expert platforms. This is a side effect of being an old and long-term toolkit. If you run into it, and wonder if it can serve your purpose, just ask an expert.

Most \TeX solutions are open source and come for free as well. Of course if you want a specific solution or want support beyond what is offered on mailing lists and forums you should be willing to pay for the hours spent. For a professional publisher (of whatever kind) this is not a problem, if only because any other solution also will cost something. It is hard to come up with a general estimate. A popular measure of typesetting costs is the price per page, which can range from a couple of euro's per page to two digit numbers. We've heard of cases where initial setup costs were charged. If not much manual intervention is needed a \TeX solution mostly concerns initial costs.

Let's return to the main question "Why use \TeX ?" in which you can replace \TeX by one of the macro packages built on top of it, for instance ConTeXt. If an (somewhat older) organization considers using \TeX it should also ask itself, why it wasn't considered long ago already? For sure there have been developments in \TeX engines (in ConTeXt we use LuaTeX) as well as possibilities of macro packages but if you look at the documents produced with them, there is not that much difference with decades ago. Processing has become faster, some things have become easier, but new technologies have always been supported as soon as they showed up. Advertising is often just repeating an old message.

The \TeX ecosystem was among the first in supporting for instance OpenType, and the community even made sure that there were free fonts available. A format like pdf was supported as soon as it showed up and \TeX was the first to demonstrate what advanced

features were there and it shows again how it is possible to adapt \TeX to changes in its environment. Processing xml using \TeX has never been a big deal and if that is a reason to look at this already old and mature technology, then an organization can wonder if years and opportunities (for instance for publishing on demand or easy updating of manuals) have been lost. Of course there are (and have been) alternative tools but the arguments for using \TeX or not are not much different now. It can be bad marketing of open and free software. It can be that \TeX has been around too long. It can also be that its message was not understood yet. On the other hand, in software development it's quite common to reinvent wheels and present old as new. It's never too late to catch on.

What's to stay, what's to go 4

4.1 Introduction

The following text was written as preparation for a 2018 talk at BachtEX, which has this theme. It's mostly a collection of thoughts. It was also more meant as a wrapup for the presentation (possibly with some discussions) than an article.

4.2 Attraction

There are those movies where some whiz-kid sits down behind a computer, keys in a few commands, and miracles happen. Ten fingers are used to generate programs that work immediately. It's no problem to bypass firewalls. There is no lag over network connections. Checking massive databases is no big deal and there's even processing power left for real time visualization or long logs to the terminal.

How boring and old fashioned must a regular edit-run-preview cycle look compared to this. If we take this 2018 movie reality as reference, in a time when one can suck a phone empty with a simple connection, pull a hard drive from a raid five array and still get all data immediately available, when we can follow realtime whoever we want using cameras spread over the country, it's pretty clear that this relatively slow page production engine TeX has no chance to survive, unless we want to impress computer illiterate friends with a log flying by on the console (which in fact is used in movies to impress as well).

On YouTube you can find these (a few hours) sessions where Jacob Collier harmonizes live in one of these Digital Audio Workstation programs. A while later on another channel June Lee will transcribe these masterpieces into complex sheets of music by ear. Or you can watch the weekly Wintergatan episodes on building the Marble Machine from wood using drilling, milling, drawing programs etc. There are impressive videos of multi-dimensional led arrays made by hand and controlled by small computers and robots that solve Rubic Cubes. You can be impressed by these Animusic videos, musicians show their craftsmanship and interesting informative movies are all over the place. I simply cannot imagine millions of kids watching a TeX style being written in a few hours. It's a real challenge for an attention span. I hope to be proven wrong but I fear that for the upcoming generation it's probably already too late because the 'whow' factor of TeX is low at first encounter. Although: picking up one of Don Knuths books can have that effect: a nice mixture of code, typesetting and subtle graphics, combined with great care, only possible with a system like TeX.

: Biology teaches us that 'cool' is not a recipe for 'survival'. Not all designs by nature look cool, and it's only efficiency and functionality that matters. Beauty sometimes matters too but many functional mechanisms can do without. So far

\TeX and its friends were quite capable to survive so there must be something in it that prevents it to be discarded. But survival is hard to explain. So far \TeX just stayed around but lack of visual attraction is a missing competitive trait.

4.3 Satisfaction

Biology also teaches us that chemistry can overload reason. When we go for short-term pleasure instead of long-term satisfaction (Google for Simon Sinek on this topic), addiction kicks in (for instance driven by crossing the dopamine thresholds too often, Google for Robert Sapolsky). Cool might relate more to pleasure while satisfaction relates to an effort. Using \TeX is not that cool and often takes an effort. But the results can be very satisfying. Where ‘cool’ is rewarding in the short term, ‘satisfaction’ is more a long term effect. So, you probably get the best (experience) out of \TeX by using it a lifetime. That’s why we see so many old \TeX ies here: many like the rewards.

If we want to draw new users we run into the problem that humans are not that good in long term visions. This means that we cannot rely on showing cool (and easy) features but must make sure that the long term reward is clear. We can try to be ‘cool’ to draw in new users, but it will not be the reason they stay. Instant success is important for kids who have to make a report for school, and a few days “getting acquainted with a program” doesn’t fit in. It’s hard to make kids addicted to \TeX (which could be a dubious objective).

: As long as the narrative of satisfaction can be told we will see new users. Meetings like Bacho \TeX is where the narrative gets told. What will happen when we no longer meet?

4.4 Survival

Survival relates to improvements, stability and discarding of weak aspects. Unfortunately that does not work out well in practice. Fully automated multi-columns typesetting with all other elements done well too (we just mention images) is hard and close to impossible for arbitrary cases, so nature would have gotten rid of it. Ligatures can be a pain especially when the language is not tagged and some kind of intelligence is needed to selectively disable them. They are the tail of the peacock: not that handy but meant to be impressive. Somehow it stayed around in automated typesetting, in biology it would be called a freak of nature: probably a goodbye in wildlife. And how about page breaks on an electronic device: getting rid of them would make the floating figures go away and remove boundary conditions often imposed. It would also make widows and clubs less of a problem. One can even wonder if with page breaks the windows and clubs are the biggest problems, and if one can simply live with them. After all, we can live with our own bodily limitations too. After all, (depending on what country you live in) you can also live with bad roads, bad weather, pollution, taxes, lack of healthcare for many, too much sugar in food, and more.

: Animals or plants that can adapt to live on a specific island might not survive elsewhere. Animals or plants introduced in an isolated environment might quickly dominate and wipe out the locals. What are the equivalents in our \TeX ecosystem?

4.5 Niches

But arguments will not help us determine if \TeX is the fittest for survival. It's not a rational thing. Humans are bad in applying statistics in their live, and looking far ahead is not a treat needed to survive. Often nature acts in retrospect. (Climbing mount probability by Richard Dawkins). So, it doesn't matter if we save time in the future if it complicates the current job. If governments and companies cannot look ahead and act accordingly, how can we extrapolate software (usage) or more specifically typesetting demands. Just look at the political developments in the country that hosts this conference. Could we have predicted the diminishing popularity of the eu (and disturbing retrograde political mess in some countries) of 2018 when we celebrated the moment Poland joining the eu at a Bachtex campfire?

Extrapolating the future quality of versions of \TeX or macro packages also doesn't matter much. With machine learning and artificial intelligence around the corner and with unavoidable new interfaces that hook into our brains, who knows what systems we need in the future. A generic flexible typesetting system is probably not the most important tool then. When we discuss quality and design it gets personal so a learning system that renders neutrally coded content into a form that suits an individual, demands a different kind of tool than we have now.

On the short term (our live span) it makes more sense to look around and see how other software (ecosystems) fare. Maybe we can predict \TeX 's future from that. Maybe we can learn from others mistakes. In the meantime we should not flatter ourselves with the idea that a near perfect typesetting system will draw attention and be used by a large audience. Factors external to the community play a too important role in this.

: It all depends on how well it fits into a niche. Sometimes survival is only possible by staying low on the radar. But just as we destroy nature and kill animals competing for space, programs get driven out of the software world. On a positive note: in a project that provides open (free) math for schools students expressed to favour a printed book over web-only (one curious argument for web was that it permits easier listening to music at the same time).

4.6 Dominance

Last year I installed a bit clever (evohome) heating control system. It's probably the only "working out of the box" system that supports 12 zones but at the same time it has a rather closed interface as any other. One can tweak a bit via a web interface but that one works by a proxy outside so there is a lock in. Such a system is a gamble because

it's closed and we're talking of a 20 year investment. I was able to add a layer of control (abusing LuaTeX as Lua engine and ConTeXt as library) so let's see. When I updated the boiler I also reconfigured some components (like valves) and was surprised how limited upgrading was supported. One ends up with lost settings and weird interference and it's because I know a bit of programming that I kept going and managed to add more control. Of course, after a few weeks I had to check a few things in the manuals, like how to enter the right menu.

So, as the original manuals are stored somewhere, one picks up the smart phone and looks for the manual on the web. I have no problem with proper pdf as a manual but why not provide a simple standard format document alongside the fancy folded A3 one. Is it because it's hard to produce different instances from one source? Is it because it takes effort? We're talking of a product that doesn't change for years.

: The availability of flexible tools for producing manuals doesn't mean that they are used as such. They don't support the survival of tools. Bad examples are a threat. Dominant species win.

4.7 Extinction

When I was writing this I happened to visit a bookshop where I always check the SciFi section for new publications. I picked out a pocket and wondered if I had the wrong glasses on. The text was wobbling and looked kind of weird. On close inspection indeed the characters were kind of randomly dancing on the baseline and looked like some 150 dpi (at most) scan. (By the way, I checked this the next time I was there by showing the book to a nephew.) I get the idea that quite some books get published first in the (more expensive) larger formats, so normally I wait till a pocket size shows up (which can take a year) so maybe here I had to do with a scan of a larger print scaled down.

What does that tell us? First of all that the publisher doesn't care about the reader: this book is just unreadable. Second, it demonstrates that the printer didn't ask for the original pdf file and then scaled down the outline copy. It really doesn't matter in this case if you use some high quality typesetting program then. It's also a waste of time to talk to such publishers about quality typesetting. The printer probably didn't bother to ask for a pdf file that could be scaled down.

: In the end most of the publishing industry will die and this is just one of the symptoms. Typesetting as we know it might fade away.

4.8 Desinterest

The newspaper that I read has a good reputation for design. But why do they need to drastically change the layout and font setup every few years? Maybe like an animal marking his or her territory a new department head also has to put a mark on the layout. Who knows. For me the paper became pretty hard to read: a too light font that suits

none of the several glasses that I have. So yes, I spend less time reading the paper. In a recent commentary about the 75 year history of the paper there was a remark about the introduction of a modern look a few decades ago by using a sans serif font. I'm not sure why sans is considered modern (most handwriting is sans) and to me some of these sans fonts look pretty old fashioned compared to a modern elegant serif (or mix).

: If marketing and fashion of the day dominate then a wrong decision can result in dying pretty fast.

4.9 Persistence

Around the turn of the century I had to replace my cd player and realized that it made more sense to invest in ripping the cd's to flac files and use a decent dac to render the sound. This is a generic approach similar to processing documents with \TeX and it looks as future proof as well. So, I installed a virtual machine running SlimServer and bought a few SlimDevices, although by that time they were already called SqueezeBoxes.

What started as an independent supplier of hardware and an open source program had gone the (nowadays rather predictable) route of a buy out by a larger company (Logitech). That company later ditched the system, even if it had a decent share of users. This "start something interesting and rely on dedicated users", then "sell yourself (to the highest bidder)" and a bit later "accept that the product gets abandoned" is where open source can fail in many aspects: loyal users are ignored and offended with the original author basically not caring about it. The only good thing is that because the software is open source there can be a follow up, but of course that requires that there are users able to program.

I have 5 small boxes and a larger transporter so my setup is for now safe from extinction. And I can run the server on any (old) linux or MS Windows distribution. For the record, when I recently connected the 20 year old Cambridge CD2 I was surprised how well it sounded on my current headphones. The only drawback was that it needs 10 minutes for the transport to warm up and get working.

In a similar fashion I can still use \TeX , even when we originally started using it with the only viable quality dvi to PostScript backend at that time (dvipsone). But I'm not so sure what I'd done if I had not been involved in the development of pdf \TeX and later Lua \TeX . As an average user I might just have dropped out. As with the cd player, maybe someone will dust off an old \TeX some day and maybe the only hurdle is to get it running on a virtual retro machine. Although . . . recently I ran into an issue with a virtual machine that didn't provide a console after a kvm host update, so I'm also getting pessimistic about that escape for older programs. (Not seldom when a library update is forced into the Lua \TeX repository we face some issue and it's not something the average user want (or is able to) cope with.)

: Sometimes it's hard to go extinct, even when commerce interfered at some point. But it does happen that users successfully take (back) control.

4.10 Freedom

If you buy a book originating in academia written and typeset by the author, there is a chance that it is produced by some flavour of \TeX and looks quite okay. This is because the author could iterate to the product she or he likes. Unfortunately the web is also a source of bad looking documents produced by \TeX . Even worse is that many authors don't even bother to set up a document layout properly, think about structure and choose a font setup that matches well. One can argue that only content matters. Fine, but then also one shouldn't claim quality simply because \TeX has been used.

I've seen examples of material meant for bachelor students that made me pretend that I am not familiar with \TeX and cannot be held responsible. Letter based layouts on A4 paper, or worse, meant for display (or e-book devices) without bothering to remove the excessive margins. Then these students are forced to use some collaborative \TeX environment, which makes them dependent on the quality standards of fellow students. No wonder that one then sees dozens of packages being loaded, abundant copy and paste and replace of already entered formulas and interesting mixtures of inline and display math, skips, kerns and whatever can help to make the result look horrible.

: Don't expect enthusiast new users when you impose \TeX but take away freedom and force folks to cooperate with those with lesser standards. It will not help quality \TeX to stay around. You cannot enforce survival, it just happens or not, probably better with no competition or with a competition so powerful that it doesn't bother with the niches. In fact, keeping a low profile might be best! The number of users is no indication of quality, although one can abuse that statistic selectively?

4.11 Diversity

Diversity in nature is enormous. There are of course niches, but in general there are multiple variants of the same. When humans started breeding stock or companion animals diversity also was a property. No one is forcing the same dog upon everyone or the same cow. However, when industrialization kicks in things become worse. Many cows in our country share the same dad. And when we look at for instance corn, tomatoes or whatever dominance is not dictated by what nature figures out best, but by what commercially makes most sense, even if that means that something can't reproduce by itself any longer.

In a similar way the diversity of methods and devices to communicate (on paper) at some point turns into commercial uniformity. The diversity is simply very small, also in typesetting. And even worse, a user even has to defend her/himself for a choice of system (even in the \TeX community). It's just against nature.

: Normally something stays around till it no longer can survive. However, we humans have a tendency to destroy and commerce is helping a hand here. In that respect it's a surprise that \TeX is still around. On the other hand, humans

also have a tendency to keep things artificially alive and even revive. Can we revive \TeX in a few hundred years given the complex code base and Make infrastructure?

4.12 Publishing

What will happen with publishing? In the production notes of some of my recently bought books the author mentions that the first prints were self-published (either or not sponsored). This means that when a publisher “takes over” (which still happens when one scales up) not much work has to be done. Basically the only thing an author needs is a distribution network. My personal experience with for instance cd’s produced by a group of musicians is that it is often hard to get it from abroad (if at all) simply because one needs a payment channel and mail costs are also relatively high.

But both demonstrate that given good facilitating options it is unlikely that publishers as we have now have not much chance of survival. Add to the argument that while in Gutenbergs time a publisher also was involved in the technology, today nothing innovative comes from publishers: the internet, ebook devices, programs, etc. all come from elsewhere. And I get the impression that even in picking up on technology publishers lag behind and mostly just react. Even arguments like added value in terms of peer review are disappearing with the internet where peer groups can take over that task. Huge amounts of money are wasted on short-term modern media. (I bet similar amounts were never spent on typesetting.)

: Publishers, publishing, publications and their public: as they are now they might not stay around. Lack of long term vision and ideas and decoupling of technology can make sure of that. Publishing will stay but anyone can publish; we only need the infrastructure. Creativity can win over greed and exploitation, small can win over big. And tools like \TeX can thrive in there, as it already does on a small scale.

4.13 Understanding

“Why do you use \TeX ?” If we limit this question to typesetting, you can think of “Why don’t you use MS Word?” “Why don’t you use Indesign?”, “Why don’t you use that macro package?”, “Why don’t you use this \TeX engine?” and alike. I’m sure that most of the readers had to answer questions like this, questions that sort of assume that you’re not happy with what you use now, or maybe even suggest that you must be stupid not to use . . .

It’s not that easy to explain why I use \TeX and/or why \TeX is good at the job. If you are in a one-to-one (or few) sessions you can demonstrate its virtues but ‘selling’ it to for instance a publisher is close to impossible because this kind of technology is rather unknown and far from the click-and-point paradigm. It’s even harder when students get accustomed to these interactive books from wherein they can even run code snippets although one can wonder how individual these are when a student has the web

as a source of solutions. Only after a long exposure to similar and maybe imperfect alternatives books will get appreciated.

For instance speaking of “automated typesetting” assumes that one knows what typesetting is and also is aware that automated has some benefits. A simple “it’s an xml to pdf converter” might work better but that assumes xml being used which for instance not always makes sense. And while hyphenation, fancy font support and proper justification might impress a \TeX user it often is less of an argument than one thinks.

The “Why don’t you” also can be heard in the \TeX community. In the worst case it’s accompanied by a “... because everybody uses ...” which of course makes no sense because you can bet that the same user will not fall for that argument when it comes to using an operating system or so. Also from outside the community there is pressure to use something else: one can find defense of minimal markup over \TeX markup or even html markup as better alternative for dissemination than for instance pdf or \TeX sources. The problem here is that old-timers can reflect on how relatively wonderful a current technique really is, given changes over time, but who wants to listen to an old-timer. Progress is needed and stimulating (which doesn’t mean that all old technology is obsolete). When I watched Endre eNerd’s “The Time Capsule” blu-ray I noticed an Ensoniq Fizmo keyboard and looked up what it was. I ended up in interesting reads where the bottom line was “Either you get it or you don’t”. Reading the threads rang a bell. As with \TeX , you cannot decide after a quick test or even a few hours if you (get the concept and) like it or not: you need days, weeks, or maybe even months, and some actually never really get it after years.

: It is good to wonder why you use some program but what gets used by others depends on understanding. If we can’t explain the benefits there is no future for \TeX . Or more exact: if it no longer provide benefits, it will just disappear. Just walk around a gallery in a science museum that deals with computers: it can be a bit pathetic experience.

Who knows ...

Introduction

How stable is ConTeXt? This question is hard to answer. For instance MkII hasn't changed for years and seems to work quite well: no changes equals stability. Those who use it can do with what it offers. The potentially sensitive dependencies on for instance fonts are probably absent because there is not much development in the 8 bit fonts arena. As long as these are available we're okay, in fact, OpenType fonts are more a moving target and therefore less stable.

What do we mean by stable? The fundamental differences between an 8 bit engine (and fonts) and an Unicode aware engine able to handle OpenType fonts is substantial which is why we dropped some functionality and added some relevant new. One can consider that a problem but in practice using fonts has become easier so no one is hurt by it. Here we need to keep in mind that pdfTeX is really stable: it uses fonts and technology that doesn't change. On the other hand XeTeX and LuaTeX follow new trends. Thereby XeTeX uses libraries, which introduces a dependency and instability, while LuaTeX assumes solutions in Lua which means that users and macro writers can tweak and thereby also introduce instability (but at least one can adapt that code).

Due to the way the user interface is set up, it is unlikely that ConTeXt will change. But the fact that we now have Lua available means that many commands have been touched. Most behave compatible, some have more functionality, and of course we have a Lua interface. We include a lot of support code which also lessens dependencies.

The user input is normally TeX but when you use xml the move to MkIV meant that we dropped the MkII way of dealing with it in favour of a completely new mechanism. I get the impression that those using xml don't regret that change. Talking of stability the MkIV xml interface is typically a mechanism that is stable and might change little. We can add new trickery but the old stays as it is.

If we look at the output, there is dvi and pdf. In MkII the dvi could become PostScript. As there are different dvi post-processors the backend code was using a plug-in model. Contrary to other macro packages there was only one so called format that could adapt itself to the required (engine specific) output. A ConTeXt run has always been managed by a wrapper so users were not bothered much by what TeX engine they used and/or what backend was triggered. This changed with MkIV where we use just LuaTeX, always produce pdf and optionally can export xml. But again the run is managed by a wrapper, which incidentally is written in Lua and thereby avoids dependencies on for instance Perl, Ruby or Python, which are moving targets, use libraries and additional user code, and thereby are potentially instable too.

The pdf code that is produced is a mix of what the engine spits out and what the macro package injects. The code is normally rather simple. This means that it's no big deal to support the so called standards. It also means that we can support advanced interactivity and other features but these also depends on the viewers used. So, stability here is more fluent, for instance because the pdf standard evolves and/or we need to adapt to viewers. Special demands like tagged pdf have been supported right from the start but how that evolves depends mostly on input from users who need it. Again, that is less important (and crucial) for stability than the rendering capabilities.

The fact that we use Lua creates a dependency on that language but the reason that we use it is *because* it is so stable. We follow the updates and so far that worked out well. Now, say that we had a frozen version of ConTeXt 2010 and LuaTeX 1.09 that uses Lua 5.3, would that work? First of all, in 2010 LuaTeX itself was evolving so the answer is probably “no”, unless one adds a few compatibility patches. I’m not going to try it. The change from 5.1 to 5.2 to 5.3 was not really a problem I think and the few issues could be dealt with easily. If you want long term stability and use a lot of Lua code you can take it into account when coding. Avoiding external libraries is a good start.

Fonts are more than before moving targets. So, if you want stability there you should save them with your document source. The processing of them has evolved and has been improved over time. By now it’s rather stable. More recent code can catch more issues and fixes are relatively easy. But it’s an area that you always need to check when you update an old distribution. The same is true for language related hyphenation patterns and script specific support. The community is no longer leading in the math department either (OpenType math is a Microsoft invention). But, the good news is that the TeX ecosystem is always fast to adapt and can also often provide more functionality.

Vertical spacing, in fact spacing in general is an area that can always be improved, so there is where you can expect changes. The same is true for side floats or mechanisms where content is somehow attached to other moving content, for instance marginal notes.

But code dealing with fonts, color, scripts, structure, and specific features that once written don’t need more, will not change that much. As mentioned for fonts, like any resource, we also depend on third parties. Colors can relate to standards, but their main properties are unchanged. Support for specific scripts can (and will) be improved due to user input and demands so there the users also influence stability. Structure doesn’t really influence the overall rendering, but the way you set it up does, but that’s user styling. Of course during the transition from MkII to MkIV and the evolution of LuaTeX things could be broken, but fixing something structural seldom relates to rendering. If for instance we improve the interpretation of bibTeX input, which can be real messy, that involves data processing, nor rendering. When we improve support for the apa standard, which is complex, it might involve rendering but then that’s asked for and expected. One cannot do better than the input permits.

Publishers

When discussing stability and especially stability as requirement we need to look at the way ConTeXt is used. So let's look at a few scenarios. Say that a publisher gets a camera ready book from an author in pdf format. In that case the author can do all tweaks needed. Now say that the publisher also wants the source code in a format that makes reuse possible.

But let's face reality. Will that publisher really reformat the document in pdf again? It's very unlikely. First of all the original pdf can be kept, and second, a reformat only makes sense after updating the content or going for a completely different layout. It's basically a new book then. In that case literal similarity of output is irrelevant. It is a cheap demand without much substance.

When the source is used for a different purpose the tool used to make the pdf is irrelevant. In that case the coding of the source can matter. If it is in some dialect of \TeX , fine, one has to convert it anyway (to suit the other usage). If there is an xml export available, fine too as it can be transformed, given that the structure is rich enough, something that is unlikely to have been checked when the original was archived. Then there could have been the demand for a document in some other format and who can guarantee stability of the tools used there? Just look at how Microsoft Word evolved, or for that matter, its competitors. On the average \TeX is more stable as one can snapshot a \TeX tree and run binaries for years, if needed, in a virtual machine.

So, I don't think that a publisher is of any relevance in the discussion about stability. Even if we can clearly define what a publisher is, I doubt if publishers themselves can be considered long term stable organizations. Not today. I'm not sure if (especially the large) publishers really deserve a place in the discussion about stability but I'm willing to discuss that when I run into one.

The main problem that an author can face when being confronted with the stability issue this way is that the times are long gone that publishers have a clue about what \TeX is, how it evolved and how it always had to and did adapt to changing requirements. If you're lucky you will run into someone who does know all this. They're normally a bit older and have seen the organization from any angles and therefore are fun to work with.

But even then, rendering issues are often not high on their agenda. Outsourcing often has become the modus operandi which basically brings us to the second group involved in this discussion: suppliers.

Suppliers

I don't know many suppliers other than the ones we ran into over a few decades. At least where I live the departments that are responsible for outsourcing typesetting like to deal with only a few large suppliers, interestingly because they assume that they

are stable. However, in my experience hardly any of those seem to have survived. (Of course one can wonder if long term commitment really is that important in a world where companies change so fast.) This is somewhat obscured by the fact that publishers themselves merge, reorganize, move people around, etc. so who can check on the stability of suppliers. It is definitely a fact that at least recently hardly any of them played a role of any relevance in the development of stable tools. In the past the membership of \TeX user groups contained people working at publishers and suppliers but that has changed.

Let's focus on the suppliers that somehow use \TeX and let's consider two kind of suppliers: small ones, one were only a few people work, and large ones. The small ones depend on stable \TeX distributions, like $\text{\TeX}Live$ where they can get the resources from: styles, fonts, patterns, binaries. If they get the authors \TeX files they need to have that access. They have to rework that input into what the customer demands and that likely involves tweaks. So, maybe they have developed their own additional code. For that code, stability is their own responsibility. Did they tweak core code of a macro package? Fine, but you might have it coming when you update. You cannot expect the evolving free meal world to stick to your commercial needs. A supplier can play safe and somehow involve the developers of macro packages or consult them occasionally, but does that really happen often? Interesting is that a few times that I was asked for input it was also wrapped in obscurity, as if some holy grail of styling was involved, while it's quite likely that the developer of a macro package can write such a style (or extra code) easily and probably also better. There really is not that much unique code around.

Small suppliers can be on mailing lists where they can contribute, get feedback, provide testing, etc. They are part of a process and as such have some influence on stability. If they charge by the page, then a change in their tools can be reflected in what they charge. Basically redoing a book (or so) after a decade is doing a new job. And adapting to some new options in a package, as part of a typesetting job is probably no big deal. Is commercial really more stable than open source free software? Probably not, except from open source software developers whose real objective is to eventually sell their stuff to some company (and cash) and even accept it to be ditched. Small suppliers are more flexible.

The large suppliers are a different group. They often guard their secrets and stay in the dark. They probably seldom share (fundamental) code and information. If they are present in a community it can be for marketing reasons. If at some point a large supplier would demand stability, then my first response would be: sure I can make you a stable setup and maybe even provide intermediate patches but put your money where your mouth is. But that never happened and I've come to the conclusion that we can safely ignore that group. The \TeX user groups create distributions and have for instance funded font development and it are the common users who paid for that, not the scale ones. To some extent this is actually good because large (software related) organizations often have special agendas that can contradict what we aim at in the long term.

From the authors perspective there is a dilemma here. When you submit to a publisher who outsources, it can be a demand to deliver in a specific \TeX format. Often a pdf comes with the source then, so that the intended rendering is known. Then that source goes to a supplier who then (quite likely) redoes a lot of the coding in some stable subset, maybe even in a very old version of the macro package. If I were such an author I'd render the document in 'as stupid as possible mode' because you gain nothing by spending time on the looks. So, stability within the package that you use is easy and translation from one to another probably also. It's best to check beforehand what will happen with your source and let stability, if mentioned, be their problem. After all they get paid for it.

Suppliers seldom know $\text{Con\TeX}t$. An interesting question is if they really know the alternatives well, apart from the bit they use. A well structured $\text{Con\TeX}t$ source (or probably any source) is often easy to convert to another format. You can assume that a supplier has tools for that (although we're often surprised about the poor quality of tools used). Often the strict demand for some kind of format is an excuse for lack of knowledge. Unfortunately you need a large author base to change that attitude.

Authors

Before we move to some variants of the above, first I will look at stability from the authors perspective. When a book is being written the typesetting more or less happens as part of the process. The way it looks can influence the way you write and vice versa. Once the book is done it can go in print and, unless you were using beta versions of $\text{Con\TeX}t$ and updated frequently. Normally you will try to work in a stable setup. Of course when a user asks for additional features while working on a project, he or she should also accept other beta features and side effects.

After a few years an author might decide to update the book. The worst that can happen is that the code doesn't run with the latest $\text{Con\TeX}t$. This is not so likely because commands are upward compatible. However, the text might come out a bit different, for instance because different fonts or patterns are used. But on the average paragraphs will come out the same in \TeX . You can encounter differences in the vertical spacing and page breaks, because that is where improvements are still possible. If you use conceptually and implementation wise complex mechanism like side floats, you can also run into compatibility issues. But all these don't really matter much because the text will be updated anyway and fine-tuning of page breaks (if at all) happens at the end. The more you try to compete with desk top publishing, and the more tweaks you apply, the greater is the risk that you introduce instability. It is okay for a one-time job, but when you come back to it after a decade, be prepared for surprises.

Even if you stick to the original coding, it makes sense to sacrifice some of that stability if new mechanisms have become available. For instance, if you use MetaPost, better ways to solve your problem might have become available. Or if your document is 15 years old, a move from MkII to MkIV is a valid option, in which case you might also consider using the latest fonts.

Of course, when you made a style where you patched core code, you can expect problems, because anything not explicitly mentioned in the interface definition files is subjected to change. But you probably see that coming anyway.

So, is an author (or stand alone user) really dependent on stability? Probably less than thought. In fact, the operating system, internet and browsers, additional tools: all change over time and one adapts. It's something one can live with. Just see how people adapt to phones, tablets, social media, electric cars, etc. As long as the document processes and reasonable output is generated it's fine. And that is always what we aim at! After all we need to be able to use it ourselves, don't we?

Projects

Although it is often overlooked as valid alternative in rendering in large scale projects, ConTeXt is perfect as component in a larger whole. Something goes in, something comes out. In a long term project one can just install a minimal distribution, write styles, and run it for ages. Use a virtual machine and we're talking decades without any change. And, when one updates, it's easy to check if all still works. Often the demands and styles are simple and predictable. It's way more likely that a hard coded solution in some large programming environment has stability issues than that the ConTeXt bit has.

If ConTeXt is used in for instance documentation of (say) software, again there is no real issue. Such documents are simple, evolve and therefore have no stable page flow, and updating ConTeXt is not needed if the once decided upon coding is stable. You don't need the latest features. We've written styles and setups for such tasks and indeed they run for ages.

It can make me smile to see how much effort sometimes goes in low quality rendering where ConTeXt could do a way better job with far less investment in time and money but where using some presumed stable toolkit is used instead, one that comes with expensive licensing, from companies that come and go but shine in marketing. (A valid question is to what extent the quality of and care for documentation reflects the core products that a company produces, at least under the hood.)

The biggest hurdle in setting up a decent efficient workflow is that it has to be seen as a project: proper analysis, proper planning, prototyping and testing, etc. You invest first and gain later. When dealing with paper many publishers still think in price per page and have problems seeing that a stable mostly automated flow in the end can result in a ridiculous low price per page, especially in typesetting on demand.

Hybrids

Last I will mention a setup that we sometimes are involved in. An author writes books and uses TeX. The publisher is okay with that and adds some quality assurance but in the end the product comes from the author. Maybe images are oursourced (not always

for the better) but these can be handled easily. It can be that a copy-editor is involved and that person also then has to use \TeX of course, or feedback to the author.

Publishers, and this really depends on knowledgeable persons, which as said can be fun to work with, can look beyond paper and also decide for additional materials, for instance web pages, interactive exercises, etc. In that case either ConTeXt input has to be available as xml (an export) or (often better) xml is the starting point for multiple output. Contrary to what is believed, there are authors out there who have no problem coding in xml directly. They think in structured content and reuse! The fact that they can hit a button in the editor and see the result in pdf helps a lot. It just works.

Here stability is either achieved by simply not updating during a project. There are however cases where an update is needed, for instance because demands changed. An example is a project where AsciiMath is used which is a moving target. Of course one can update just that module, and often that works, but not when a module uses some new improved core helpers. Another example is additional proofing options.

The budget of such projects seldom permit patching an existing distribution, so we then just update to the latest but not after checking if the used style works okay. There is no author involvement in this. Depending on the workflow, it can even be that the final rendering which involves fine tuning (side) float placement or page breaks (often educational documents have special demands) is done by us using special directives.

Such hybrid workflows are quite convenient for all parties. The publisher works with the author who likes using these tools, the author can do her or his thing in the preferred way, and we do what we're best in: supporting this. And it scales up pretty well too if needed, without much costs for the publishers.

Conclusion

So what can we conclude with respect to the demand for stability? First of course that it's important that our files keep running well. So, functionality should be stable. Freezing a distribution will make sure that during project you don't run into issues. Many ConTeXt users update frequently in order to benefit from the latest additions. Most will not be harmed by this, but when something really breaks it's users like those on the ConTeXt support list (who often also contribute in helping out other users) that are listened to first. Publishers demands play no role in this, if only because they also play no role in typesetting, and if they want to they should also contribute. The same is true for large suppliers. We're talking of free software often written without any compensation so these parties have no say in the matter unless they pay for it. It's small suppliers, authors and general users that matter most. If ConTeXt is part of a workflow that we support, of course stability is guaranteed quite well, and those paying for that never have an issue with better solutions popping up. In fact, ConTeXt is often just a tool then, one that does the job and questions about stability don't matter much in practice, as long as it does the job well.

The main engine we use, `LuaTeX`, will be quite stable from version 1.10 and we'll try to make sure that newer versions are capable of running an older `ConTeXt`, which is easier when no fundamental changes happen in the engine. Maybe a stripped down version of `LuaTeX` for `ConTeXt` can facilitate that objective even more.

Users themselves can try to stick to standard `ConTeXt` features. The more tricks you apply, the less stable your future might be. Most mechanism are not evolving but some, like those that deal with columns, might become better over time. But typesetting in columns is often a one-shot adventure anyway (and who needs columns in the future).

Of one thing users can be sure. There will never be a `ConTeXt` professional or `ConTeXt` enterprise. There is only one variant. All users get the same functionality and policies don't change suddenly. There will be no lock in to some cloud or web based service either. Of course one can hire us for support of any kind but that's independent of the distributed package. There is support by users for users on mailing lists and other media. That itself can also guard stability.

But, always keep in mind that stability and progress, either of not driven by the environment that we operate in, can be in conflict.

6.1 Introduction

Here I will shortly wrap up the state of LuaTeX and ConTeXt in fall 2018. I made the first draft of this article as preparation for the ConTeXt meeting where we also discussed the future. I updated the text afterwards to match the decisions made there. It's also a personal summary of thoughts and discussions with team members about where to move next.

6.2 The state of affairs

After a dozen years the development of LuaTeX has reached a state where adding more functionality and/or opening up more of the internals makes not much sense. Apart from fixes and maybe some minor extensions, version 1.10 is what you get. Users can do enough in Lua and there is not much to gain in convenience and performance. Of course some of the code can and will be cleaned up, as we still see the effects of going from Pascal to cweb to C. In the process consistency is on the radar so we might occasionally add a helper. But we also don't want to move too far away from the original code, which is for instance why we keep names, keys and other properties found in original TeX, which in turn leads to some inconsistencies with extensions added over time. We have to accept that.

Because LuaTeX development is closely related to ConTeXt development, especially MkIV, we've also reached the moment that we can get rid of some older code and assume the latest LuaTeX to be used. Because we do so much in Lua the question is always to what extent the benefits outweigh the drawbacks. Just in case you wonder why we use Lua extensively, the main reason is that it is easier and more efficient to manage data in this language and modern typesetting needs much data. It also permits us to extend regular TeX functionality. But, one should not overrate the impact: we still let TeX do what TeX is best at!

Performance is quite important. It doesn't make sense to create a powerful typesetting system where processing a page takes a second. We have discussed performance before since one of the complaints about LuaTeX is that it is slow. A simple, basic test is this:

```
\starttext
  \dorecurse{1000}{\input tufte \par}
\stoptext
```

This involves 1000 times loading a file (and reporting that on the console, which can influence runtime), typesetting paragraphs, splitting of a page and of course loading

fonts and saving to the pdf file. When I run this on a modest machine, I get these (relative) timings for the (about) 225 pages:

TeX engine used	pdfTeX	LuaTeX	LuajitTeX	XeTeX
runtime in seconds	2.0	3.9	3.0	8.4

Now, as expected the 8 bit pdfTeX is the winner here but LuaTeX is not doing that bad. I don't know why XeTeX is so much slower, maybe because its 64 bit binary is less optimal. I once noticed that a 64 bit pdfTeX performed worse on such a test than LuaTeX, for which I always use 64 bit binaries.

If you consider that often much more is done than in this example, you can take my word that LuaTeX quickly outpaces pdfTeX on more complex tasks. In that sense it is now our benchmark. It must be said that the MkIV code is probably a bit more efficient than the MkII code but that doesn't matter much in this simple test because hardly any macro magic happens here; it mostly tests basic font processing, paragraph building and page construction. I don't think that I can squeeze out more pages per second, at least not without users telling me where they encounter bottlenecks that don't result from their style coding. It's no problem to write inefficient macros (or styles) so normally a user should first carefully check her/his own work. Using a more modern cpu with proper caching and an ssd helps too.

So, to summarize, we can say that with version 1.10 LuaTeX is sort of finished. Our mission is now to make LuaTeX robust and stable. Things can be added and improved, but these are small and mostly consistency related.

6.3 More in Lua

Till now I always managed to add functionality to ConTeXt without hampering performance too much. Of course the biggest challenge is always in handling fonts and common features like color because that all happens in Lua. So, the question is, what if we delegate more of the core functionality to Lua? I will discuss a few options because the ConTeXt developers and users need to agree on the path to follow. One question there is, are the possible performance hits (which can be an inconvenience) compensated by better and easier typesetting.

Fonts, colors, special typesetting features like spaced kerning, protrusion, expansion, but also dropped caps, line numbering, marginal notes, tables, structure related things, floats and spacing are not open for much discussion. All the things that happen in Lua combined with macros is there and will stay. But how about hyphenation, paragraph building and page building? And how about a leaner and meaner, future safe engine?

Hyphenation is handled in the TeX core. But in ConTeXt already for years one can also use a Lua based variant. There is room for extensions and improvements there. Interesting is that performance is more or less the same, so this is an area where we might switch to the Lua method eventually. It compares to fonts, where node mode is more or less the standard and base mode the old way.

Building the paragraphs in Lua is also available in MkIV, although it needs an update. Again performance is not that bad, so when we add features not possible (or hard to do) in regular \TeX , it might actually pay off to default to the par builder written in Lua.

The page builder is also doable in Lua but so far I only played a bit with a Lua based variant. I might pick up that thread. However, when we would switch to Lua there, it might have a bit of a penalty, unless we combine it with some other mechanisms which is not entirely trivial, as it would mean a diversion from the way \TeX does it normally.

How about math? We could at some point do math rendering in Lua but because the core mechanism is the standard, it doesn't really makes much sense. It would also touch the soul of \TeX . But, I might give it a try, just for fun, so that I can play with it a bit. It's typically something for cold and rainy days with some music in the background.

We already use Lua in the frontend: locating and reading files in \TeX , xml, Lua and whatever input format. Normalization and manipulation is all active and available. The backend is also depending on Lua, like support for special pdf features and exporting to xml. The engine still handles the page stream conversion, font inclusion and object management.

The inclusion of images is also handled by the engine, although in ConTeXt we can delegate pdf inclusion to Lua. Interesting is that this has no performance hit.

With some juggling the page stream conversion can also be done in Lua, and I might move that code into the ConTeXt distribution. Here we do have a performance hit: about one second more runtime on the 14 seconds needed for the 300 page LuaTeX manual and just over more than half a second on a 11 second LuajitTeX run. The manual has lots of tables, verbatim, indices and uses color as well as a more than average number of fonts and much time is spent in Lua. So there is a price to pay there. I tried to speed that up but there is not much to gain there.

So, say that we default to Lua based hyphenation, which enables some new functionality, Lua based par building, which permits some heuristics for corner cases, and Lua based page building, which might result in more control over tricky cases. A total performance hit of some 5% is probably acceptable, especially because by that time I might have replaced my laptop and won't notice the degrade. This still fits in the normal progress and doesn't really demand a roadmap or wider acceptance. And of course we would still use the same strategies as implemented in traditional \TeX as default anyway.

6.4 A more drastic move

More fundamental is the question whether we delegate more backend activity to Lua code. If we decide to handle the page stream in Lua, then the next question is, why not also delegate object management and font inclusion to Lua. Now, keep in mind that this is all very ConTeXt specific! Already for more than a decade we delegate a lot to Lua,

and also we have a rather tight control over this core functionality. This would mean that ConTeXt doesn't really need the backend code in the engine.⁴

That situation is actually not unique. For instance, already for a while we don't need the LuaTeX font loader either, as loading the OpenType files is done in Lua. So, we could also get rid of the font loader code. Currently some code is shared with the font inclusion in the backend but that can be isolated.

You can see a \TeX engine as being made from several parts, but the core really concerns only two processes: reading, storing and expanding macros on the one hand, and converting a stream of characters into lines, paragraphs, pages etc. Fonts are mostly an abstraction: they are visible in so called glyph nodes as font identifier (a number) and character code (also a number) properties. The result, nowadays being pdf, is also an abstraction: at some point the engine converts the to be shipped out box in pdf instructions, and in our case, relatively simple ones. The backend registers which characters and fonts are used and also includes the right resources. But, the backend is not part of the core as such! It has been introduced in pdfTeX and is a so called extension.

So, what does that all mean for a future version of ConTeXt and LuaTeX? It means that we can decide to follow up with a ConTeXt that does more in Lua, which means not hard coded in a binary, on the one hand, but that we can also decide to strip the engine from non-core code. But, given that LuaTeX is also used in other macro packages, this would mean a different engine. We cannot say that LuaTeX is stable when we also experiment with core components.

We've seen folks picking up experimental versions assuming that it is a precursor to official code. So, in order to move on we need to avoid confusion: we need to use another name. Choosing a name is always tricky but as Taco already registered the MetaTeX domain, and because in the ConTeXt distribution you will find references to MetaTeX, we will use that name for the future engine. Adding Lua to that name makes sense but then the name would become too long.

The main difference between MetaTeX and LuaTeX would be that the former has no file lookup library, no hardcoded font loader, and no backend generator (but possibly some helpers, and these need time to evolve). We're basically back where \TeX started but instead of coding these extensions in Pascal or C we use Lua. We're also kind of back to when we first started experimenting with LuaTeX in ConTeXt where test, write and rewrite were going in parallel. But, as said, we cannot impose that on a wide audience.

If we go for such a lean and mean follow up, then we can also do a more drastic cleanup of obsolete code in ConTeXt (dating from e-\TeX, pdfTeX, Aleph, etc.). We then are sort of back to where it all started: we go back to the basics. This might mean dropping some primitives (one can define them as dummy). Of course we could generalize some of the ConTeXt code to provide the kicked out functionality but would that pay off? Probably not.

⁴ For generic packages like TikZ we (can) provide some primitive emulators, which is rather trivial to implement.

Just for the record: replacing the handling of macros, registers, grouping, etc. to Lua is not really an option as the performance hit would make a large system like ConTeXt sort of unusable: it's no option and not even considered (although I must admit that I have some experimental Lua based \TeX parser code around).

It is quite likely that building MetaTeX from source for the moment will be an option to the build script. But we can also decide to simplify that process, which is possible because we only need one binary. But in general we can assume that one can generate MetaTeX and LuaTeX from the same source. A first step probably is a further isolation of the backend code. The fontloader and file handling code already can be made optional.

Given that we only need one binary (it being LuaTeX or MetaTeX) and nowadays only use OpenType fonts, one can even start thinking of a mini distribution, possibly with a zipped resource tree, something we experimented with in the early days of LuaTeX.

Another thought I have been playing with is a better separation between low level and high level ConTeXt commands, and whether the low level layer should be more generic in nature (so that one can run specific packages on top of it instead of the whole of ConTeXt) but that might not be worth the trouble.

6.5 Interlude

If we look at the future, it's good to also look at the past. Opening up \TeX the way we did has many advantages but also potential drawbacks. It works quite well in ConTeXt because we ship an integrated package. I don't think that there are many users who kick in their own callbacks. It is possible but completely up to the user to make sure things work out well. Performance hits, interference, crashes: those who interfere with the internals can sort that out themselves. I'm not sure how well that works out in other macro packages but it is a time bomb if users start doing that. Of course the documented interfaces to use Lua in ConTeXt are supported. So far I think we're not yet bitten in the tail. We keep this aspect out of the discussion.

Another important aspect is stability of the engine. Sometimes we get suggestions for changes or patches that work for a specific case but for sure will have side effects on ConTeXt. Just as we don't test L^AT_EX side effects, L^AT_EX users don't check ConTeXt. And we're not even talking of users who expect their code to keep working. A tight control over the source is important but cannot be we will not be around for ever. This means that at some point LuaTeX should not be changed any more, even when we observe side effects we want to get rid of, because these side effects can be in use. This is another argument for a stripped down engine. The less there is to mess with, the less the mess.

6.6 Audience

So how about ConTeXt itself? Of course we can make it better. We can add more examples and more documentation. We can try to improve support. The main question

for us (as developers) is who actually is our audience. From the mails coming to the ConTeXt support list it looks like a rather diverse group of users.

At TeX meetings there are often discussions about promoting TeX. I can agree on the fact that even for simple documents it makes a lot of sense to use TeX, but who will take the first hurdles? How many people really produce a lot of documents? And how many need TeX after maybe a short period of (enforced) usage at the university?

It's not trivial to recognize the possibilities and power of the LuaTeX-ConTeXt combination. We never got any serious requests for support from large organizations. In fact, we do use this combination in a few projects for educational publishers, but there it's actually the authors and editors doing the work. It's seldom company policy to use tools that efficiently automate typesetting. I dare to say that publishers are not really an audience at all: they normally delegate the task. They might accept TeX documents but let them rekey or adapt far-far-away and as cheap as possible. Thinking of it, the main reason for Don Knuth for writing TeX in the first place was the ability to control the look and feel and quality. It were developments at typesetters and publishers that triggered development of TeX. It was user demand. And the success of TeX was largely due to the unique personality and competence of the author.

System integrators qualify as audience but I fear that TeX is not considered hip and modern. It doesn't seem to matter if you can demonstrate that it can do a wonderful job efficiently and relatively cheap. Also the fact that an installation can be very stable on the long run is of no importance. Maybe that audience (market place) is all about "The more we have to program and update regularly, the merrier.". Marketing TeX is difficult.

Those who render multiple products, maintain manuals, have to render many documents automatically qualify as audience. But often company policies, preferred suppliers, so called standard tools etc. are used as argument against TeX. It's a missed opportunity.

One needs a certain mindset to recognize the potential and the question is, how do we reach that audience. Drawing a roadmap for that is not easy but worth discussing. We're open for suggestions.

6.7 Conclusion

At the ConTeXt user meeting those present agreed that moving forward this way makes sense. This means that we will explore a lean and mean MetaTeX alongside LuaTeX. There is no rush and it's all volunteer work so we will take our time for this. It boils down to some reshuffling of code so that we can remove the built-in font loader, file handling, and probably also SyncTeX because we can emulate that. Then the backend with its font inclusion code will be cleaned up a bit (we even discussed only supporting modern wide fonts). It's no big deal to adapt ConTeXt to this (so it can and will support both LuaTeX and MetaTeX). Eventually the backend might go away but now we're talking years ahead. By then we can also explore the option to make MetaTeX start out as

a Lua function call (the main control loop) and become reentrant. There will probably not be many changes to the opened up \TeX kernel, but we might extend the MetaPost part a bit (some of that was discussed at the meeting) especially because it is a nice tool to visualize big data.

As with LuaTeX development we will go in small steps so that we keep a working system. Of course LuaTeX is always there as stable fallback. The experiments will mostly happen in the experimental branch and binaries will be generated using the compile farm on the ConTeXt garden, just as happens now. This also limits testing and exploring to the ConTeXt community so that there are no side effects for mainstream LuaTeX usage.

Nowadays, instead of roadmaps, we tend to use navigational gadgets that adapt themselves to the situation. On the road by car this can mean a detour and when walking around it can be going to suggested points of interest. During the excursion at the meeting, we noticed that after the drivers (navigators) synchronized their gadget with Jano, the routes that were followed differed a bit. We saw cars in front of going a different direction and cars behind us arriving from a different direction. So, even when we talk about roadmaps, our route can be adapted to the situation.

Now here is something to think about. If you look at the \TeX community you will notice that it's an aging community. User groups seem to loose members, although the ConTeXt group is currently still growing. Fortunately we see a new generation taking interest and the ConTeXt users are a pleasant mix and it makes me stay around. I see it as an 'old timers' responsibility to have \TeX and its environment in a healthy state by the time I retire from it (although I have no plans in that direction). In parallel to the upcoming development I think we will also see a change in \TeX use and usage. This aspect was also discussed at the meeting and for sure will get a follow up on the mailing lists and future meetings. It might as well influence the decisions we make the upcoming years. So far \TeX has never failed us in it's flexibility and capacity to adapt, so let's end on that positive note.

What's in a name

7

Hans Hagen
Hasselt NL
May 2019

7.1 \TeX

I sometimes wonder how much the fact that English is the language mostly used in programming environments influences the way one looks at a program. For instance, translating the names of an operating system ‘windows’, an image manipulation program ‘photoshop’ or a text editing program ‘wordperfect’ to Dutch makes them sound kind of silly to me. The name can influence what you buy or are willing to use. These are examples of commercial programs but there are plenty examples of such naming in the open source universe too. I write this in my own bad English so that other non-English speakers can try to do a similar exercise.

So, when I was reading an article about cpu technology called ‘thread-ripper’ and after a while also saw the usual talk of yet another bunch of technologies marked as ‘stack’ and translated that to Dutch it again made me feel somewhat puzzled about such names. From there it was a small step to wondering about programming languages, and especially the ones I use: \TeX , MetaPost, and Lua.

One can even wonder to what extent the quality of programming is influenced by the names of commands and keywords. A language name ‘Basic’ sounds less serious than ‘C’. A meaningless ‘Lua’ sounds different than ‘Python’. Does using your native tongue make a difference? In Dutch and German words tend to get long. When I look at my French dictionary it is rather thin, but we might need accented characters. Words in a language like Polish can differ per usage. What if German or Spanish had been chosen as the language for what is now the United States? How would we perceive programming and what would look natural to us?

7.2 \TeX

The \TeX language comes with a lot of so called primitives built in. Many of these relate to concepts in the program. For instance, a movement in horizontal or vertical direction that can stretch or shrink depending on what the boundary conditions demand, is called ‘glue’. When discussing this in Dutch the word ‘lijm’ can be used and after seeing it a few times it might sound ok. We can probably use ‘elastiek’ (‘elastic’).

This internal concept is actually represented to the user via the interface name ‘skip’, take: `\abovedisplayskip` and `\belowdisplayskip`. Here the word ‘display’ refers

to math that gets vertical space around it and is normally typeset in a somewhat larger way compared to ‘inline’. The word ‘skip’ can be translated to ‘sprong’ (translated back we could as well get ‘jump’). But how to translate ‘display’? An internet translation can be ‘tentoonstelling’ but apart from it being a long word it sounds pretty weird for something math. The combined translation of such a command will not work well I think so probably complete different words has to be made up to describe these quantities. Taco suggested that `\bovenuitstallingkortesprong` might work for `\abovedisplayshortskip` but luckily no ordinary \TeX user will not set such parameters in a document source.

In ConTeXt we use the somewhat typographical term `wit` or `witruimte` for vertical spacing. Some parameters like `\baselineskip` can be translated directly to the Dutch `\regelafstand` which is a proper typographical term (\TeX has no concept of line height). Okay, it can become messy when we translate `\lineskip` by `\interlinespace` as that could be seen as the baseline skip too (‘interline’ comes to mind). Quite a mess. In many cases we probably would not handle the `skip` part in parameter: `leftskip` could become `\linkermarge` and `\parfillskip` can become `\paragraafuitvulling`.

Another concept is that of ‘penalty’, or in Dutch `boete`. It’s probably harder to get the combinations right, simply because they have no typographical meanings, they’re more process controllers. I fear that most translations would sound pretty weird to me. So, how do they sound to a native English speaker? Words like ‘club’ or ‘widow’ can be translated to their Dutch gender neutral counterparts ‘wees’ and ‘weduw’ but how strange does `weduwboete` sound?

The counter variables are easier. When they end on `char` that can become `karakter`. However, translating `\escapechar` with `\ontsnappingskarakter` might look a bit weird, but as that one is used very seldom, a weird one doesn’t matter much. Operators like `\advance` and `\multiply` can become `\verhoog` and `\vermenigvuldig` which doesn’t sound that strange in this context.

There are ‘rule’s and ‘box’es. The first one can be translated to ‘lijn’ which sounds quite good. But what to do with the second one. We can use ‘blok’ (which translates back to ‘block’) which is good when we start stacking things, but also with ‘doos’ which is more literal but sounds to me somewhat silly: `\hdoos{whatever}`. I’m not so sure if I would have seen that in a book about \TeX , I’d looked further into the language. The optional keywords ‘width’ etc. can be translated well into ‘breedte’ etc., so no problem there.

There are all kinds of very peculiar aspects that need a translation. For instance the (for new users intimidating) primitive `\futurelet`. The ‘future’ part is no problem as ‘toekomst’ isn’t that weird but the ‘let’ will for sure become something very long in Dutch, so we end up with `\toekomstigetoekenning`, but seeing that long one, we can consider `kijkvooruit` as reasonable alternative. It definitely leads to more verbose programming.

Expansion is a tricky one. I have no clue what would make nice translations of the primitives `\noexpand` and `\expandafter`. The Dutch ‘uitbreiden’ simply is not sounding good here. Taco Hoekwater came up with a good alternative ‘uitvouwen’ for ‘expand’ and I like that one because we then can let bookmaker (a somewhat dubious term in itself) Willi Egger organize a workshop in unfolding (instead of folding).

Talking of ‘macros’ is less a problem because there is no Dutch word for it. There are more words with no real translations: `\kern` for instance probably would need some thinking but there might be a typographical equivalent that can be used.

The e-TeX and LuaTeX extensions introduce new names, like `\detokenize`, `\boundary` and `\attribute`. The first one is hard to translate because again it relates to an internal concept: `tokens`. I get the feeling that translating each occurrence of `token` by `teken` kind of makes everything look less serious. To strip something from its special meaning, which is actually what `\detokenize` does can give weird translations: `\onttekenen` is not really a Dutch word so a complete different one has to be found that describes what happens, like `\ontwaarderen`. On the other hand, `\boundary` and `\attribute` can translate directly into `\grens` and `\attribuut` where the last one sounds mostly okay.

Just to get you thinking: how would you translate `\looseness` (`losheid`, related to linebreaking), `\deadcycles` (`zinlozelus`, in the perspective of building pages), `\pretolerance` (again line break related, here we can use something `tolerantie`) and `\prevgraf` (which is actually even in English a weird one but hardly used anyway, so Taco likes `voorloopregels`)? The easy ones are `\omit`, `\meaning`, `\number`, maybe even `\mark`. The for users often difficult to grasp ‘catcode’ can be simplified to ‘code’ which is proper Dutch. Concepts like ‘align’ translate well to ‘uitlijnen’. Short ones like `\wd` could be a problem but any two letter combination can look bad, so `\br` could do. In the same fashion `\def` is ok as it is also the start of the Dutch ‘definitie’. Mathematical terms like ‘text’, ‘script’ and ‘scriptscript’ can be confusing: ‘tekst’ will do but ‘schrift’ is strange.

Conditionals are not the hardest part: `\if` becomes `\als`, `\else` becomes `\anders` and `\or` is `\of`. However, turning `\ifcase` into `\inhetgevaldat` can be over the top. The `\every...` register variables can also be translated quite well, by using the `\elk` or `\elke` prefix. They are seldom seen at the user level so no real problem there.

The ‘group’ related commands are easy as ‘groep’ is a good Dutch equivalent. Even ‘global’ operations translate well (`globaal`). A dubious one is `\font` because we can use `\lettertype` but it’s not really a translation. The internet translations tend towards ‘fountain’ kind of things.

The concept of ‘discretionary’ again needs a decent typographical translation although `\hyphenation` can become `\afbreking`, translating `\discretionary` needs some imagination. The concept of ‘leaders’ is again something that can best be bound to something more typographical because `\leaders` turned `\leidinggevende` is not an option nor is `\leiders`.

The prefix `\un` as used in `\unhbox` can become `\ont` so that we get `\ontdoos` but I get the feeling that this one can be source of jokes. The more verbose `\pakdoosuit` (equivalent to `\unpackbox`) would do better. To translate `\unvcopy` into the gibberish `\ontdoosdecopie` is simply ridiculous and `\copieeruitgepakteodoos` is a bit long. The `\lower` and `\raise` on the other hand translate well to `\verlaag` and `\verhoog`. Keeping `\relax` untranslated sounds ok to me, because `\ontspan` really makes a language silly.

7.3 MetaPost

The \TeX language is driving a macro system while \LaTeX is a procedural language. The MetaPost language sits somewhere in between. It is still expanding all along but it looks a bit more like a programming language with its loops, assignments, conditionals, expressions and (sort of) functions. As a consequence some of what I mentioned in the previous section applies here.

Translation of for instance `truecorners` into `echtehoeken` can give the language a bit less serious image. Words like `linejoin`, `linecap` and `miterlimit` relate directly to the PostScript language so translating them also relates to translated PostScript.

The `primary`, `secondary` keywords can be nicely translated into serious counterparts `primaire` and `secundaire` which are words that are not really of Dutch origin anyway. The `precontrol` and `postcontrol` words relate to concepts but even there the verbose `controlepuntvoor` and `controlepuntachter` could do. However `punt` as translation for `point` can be confusing because we also use that for `period`. Translating `controls` and `curl` needs some imagination. Words like `tension` becoming `spanning` is still acceptable soundwise. However:

```
voor i=1 stap 2 tot 10:  
....  
eindvanvoor; % or: eindvoor
```

Kind of interesting is translating `if` into `als` because `fi` then becomes `sla` which is 'lettuce' or, when see as verb, 'hit'. The `true` and `false` keywords becoming `goed` and `fout` is no problem.

Turning `atleast` into `opzijnminst` at first sight looks strange but actually I can appreciate that one. And `tussendoortje` as translation of `interim`, I can live with that one too as it sounds funny. Concepts like 'suffixes' need thinking but `uitdr(ukking)` or more literal `expr(essie)` for `expr(ession)` are okay. The `expandafter`, `scantokens` and similar keywords share the problem with \TeX that they relate to concepts that are hard to translate.

The `redpart` and similar keywords could be translated into `rooddeel` but `roodkanaal` (meaning `redchannel`) might be better or maybe `rodecomponent`. As with \TeX grouping related keywords are no problem.

A `pencircle` becomes `pencirkel`, `odd` becomes `oneven`, `reverse` becomes `omgekeerd` (or `andersom` or `tegengesteld`). For `length` we use `lengte`, and so on. All these sound professional enough, just like ‘corner’ related keywords becoming ‘hoek’, although there a clash with ‘angle’ is possible. I’m less sure about `clipped` becoming `afgeknepen` or `begrensd` but `bounded` then needs some thinking as these all are more or less the same. The concept of ‘stroke’ maps onto ‘tekenen’ or ‘vegen’ but lucky us that one is not really used, contrary to `draw` that can map onto `teken`, while `fill` and `vul` match well too I guess.

The transformations are no problem but I’d use a directive instead: `rotated` or `roteer`, `slanted` or `schuin`, `scaled` or `schaal`, and `transform` or `transformeer`. As you can see, these have a reasonable word length too.

The concept of a `picture` is known in Dutch as `plaatje` or `tekening`: not an easy choice. Using `kleur` for `color` is no problem at all. A coordinate `pair` becomes a `paar`: close enough not to give subjective side effects. The `inner` and `outer` keywords translate well to `binnen` and `buiten` but in code it might look a bit strange.

So, in general, the translated commands are not that weird but still a graphic defined in Dutch keywords instead of English to me might look less serious.

7.4 Lua

We now arrived at a more traditional programming language. The Lua language only has a few keywords. I suppose that it’s just a matter of time before one gets accustomed to `als ... dan ... anders ... eind` instead of `if ... then ... else ... end`. The loops also translate rather well: `zolang ... doe`, `herhaal .. totdat`, `voor ... in` ... `doe ... einde` are all not that verbose. Also, with proper syntax highlighting they stand out and become abstract words. But because examples for kids are normally in Dutch, using a Dutch programming language might give a toy language feeling.

The `local` directive is a bit of a problem because it should be `lokale variabele` in order to sound ok in a sentence. The `goto` should become `ganaar` which is also two words with no space in between. The `function` keyword can become `functie`. A `coroutine` is a challenge (also conceptually); we do have `routine` but how about the `co` part?

Because Lua is such a clean language it doesn’t really end up bad. In C there are some more issues due to the abbreviated `struct`, `int`, `char`, `enum` and `typedef`. A literal translation of `void` to `leegte` to me sounds a bit strange. What to do with `unsigned`? Coming up with something (short) Dutch for `return` is not easy either. Translating `switch` into `schakelaar` looks like a bad idea but after consulting Taco using `keuze` came up. The `break` then can be `klaar` which roundtrips to ‘finished’ and `default` can be `anders` which roundtrips to ‘otherwise’ which is indeed what some languages provide.. But, there are programming languages out there that have plenty keywords and that are more challenging. But as I’m a happy Lua user I don’t have to worry about them.

7.5 Conclusion

Looking at a program source in Dutch the general feeling probably will be different. A low level bit of \TeX is the worst. For MetaPost it's bearable and for Lua it is kind of okay. But in all cases, I'm not convinced that it would give me the same feeling. The abstraction of the language due to it not being my native tongue makes a difference. This problem is not much different than what we have with popular music and songs: for non-native speakers it's basically sounds, but for a native speaker it is more clear when nonsense is sung. The same can happen to me with movies, where watching some scandinavian series is different from watching a Dutch one. In the last case one picks up different nuances, not necessarily for the best. But it can be worse: post synchronized (audio) translations can be pretty unbearable and might compare well to programs translated to for instance Dutch. So let's not discuss the way Germans would deal with this.

7.6 Side notes

We now see monospaced fonts showing up that provide ligatures for e.g. \textless= and I've seen examples where ligatures kicked in for $\text{\texttt{fi}}$. One can wonder about that but ligatures are definitely something to keep in mind when translating.

The ConTeXt macro package is normally used with the English user interface. But the design is such that one can provide different ones too; after all it started out Dutch. It is beyond the scope of this musing to discuss the problems with translating typographical concepts between languages, especially when there are no distinctive words. But it can (and has) been done.

About what ConTeXt isn't 8

8.1 Introduction

It really puzzles me why, when someone someplace asks if ConTeXt is suitable for her or is his needs, there are answers like: "You need to think of ConTeXt as being kind of plain \TeX : you have to define everything yourself." That answer probably stems from the fact that for \LaTeX you load some style that defines a lot, which you then might need to undefine or redefine, but that's not part of the answer.

In the following sections I will go into a bit more detail of what plain \TeX is and how it influences macro packages, especially ConTeXt. I'm sure I have discussed this before so consider this another go at it.

The `plain.tex` file starts with the line:

```
% This is the plain TeX format that's described in The TeXbook.
```

A few lines later we read:

```
% And don't modify the file under any circumstances.
```

So, this format is related to the \TeX reference. It serves as a template for what is called a macro package. Here I will not go into the details of macro programming but an occasional snippet of code can be illustrative.

8.2 Getting started

The first code we see in the plain file is:

```
\catcode`{\=1 % left brace is begin-group character
\catcode`{\}=2 % right brace is end-group character
\catcode`{\$=3 % dollar sign is math shift
\catcode`{\&=4 % ampersand is alignment tab
\catcode`{\#=6 % hash mark is macro parameter character
\catcode`{\^=7 \catcode`{\^{\K=7 % circumflex and uparrow are for superscripts
\catcode`{\_=8 \catcode`{\^{\A=8 % underline and downarrow are for subscripts
\catcode`{\^{\I=10 % ascii tab is a blank space
\chardef\active=13 \catcode`{\~=\active % tilde is active
\catcode`{\^{\L=\active \outer\def^{\L{\par} % ascii form-feed is "\outer\par"
```

Assigning catcodes to the braces and hash are needed in order to make it possible to define macros. The dollar is set to enter math mode and the ampersand becomes a

separator in tables. The superscript and subscript also relate to math. Nothing demands these bindings but they are widely accepted. In this respect ConTeXt is indeed like plain.

The tab is made equivalent to a space and a tilde is made active which means that later on we need to give it some meaning. It is quite normal to make that an unbreakable space, and one with the width of a digit when we're doing tables. Now, nothing demands that we have to assume ascii input but for practical reasons the formfeed character is made equivalent to a `\par`.

Now what do these `^^K` and similar triplets represent? The `^^A` represents character zero and normally all these control characters below decimal 32 (space) are special. The `^^I` is the ascii tab character, and `^^L` the formfeed. But, the ones referred to as uparrow and downarrow in the comments have only meaning on certain keyboards. So these are typical definitions that only made sense for Donald Knuth at that time and are not relevant in other macro packages that aim at standardized input media.

```
% We had to define the \catcodes right away, before the message line, since
% \message uses the { and } characters. When INITEX (the TeX initializer) starts
% up, it has defined the following \catcode values:
%
% \catcode`^^@=9 % ascii null is ignored
% \catcode`^^M=5 % ascii return is end-line
% \catcode`\\=0 % backslash is TeX escape character
% \catcode`%=14 % percent sign is comment character
% \catcode`\ =10 % ascii space is blank space
% \catcode`^^?=15 % ascii delete is invalid
% \catcode`A=11 ... \catcode`Z=11 % uppercase letters
% \catcode`a=11 ... \catcode`z=11 % lowercase letters
% all others are type 12 (other)
```

The comments above speak for themselves. Changing catcodes is one way to adapt interpretation. For instance, in verbatim mode most catcodes can best be made letter or other. In ConTeXt we always had so called catcode regimes: for defining macros, for normal text, for XML, for verbatim, etc. In MkIV this mechanism was adapted to the new catcode table mechanism available in that engine. It was one of the first things we added to LuaTeX. So, again, although we follow some standards (expectations) ConTeXt differs from plain.

```
% We make @ signs act like letters, temporarily, to avoid conflict between user
% names and internal control sequences of plain format.
```

```
\catcode`@=11
```

In ConTeXt we went a step further and when defining macros also adapted the catcode of ! and ? and later in MkIV _ . When we're in unprotected mode this applies. In addition to regular text input math is dealt with:

```
% INITEX sets up \mathcode x=x, for x=0..255, except that
```

```

%
% \mathcode x=x+"7100, for x = `A to `Z and `a to `z;
% \mathcode x=x+"7000, for x = `0 to `9.

% The following changes define internal codes as recommended in Appendix C of
% The TeXbook:

\mathcode`\^="2201 % \cdot
\mathcode`\^="3223 % \downarrow
\mathcode`\^="010B % \alpha
\mathcode`\^="010C % \beta
.....
\mathcode`\|="026A
\mathcode`\}="5267
\mathcode`\^="1273 % \smallint

```

Here we see another set of definitions but the alphabetic ones are not defined in ConTeXt, they are again bindings to the authors special keyboard.

```

% INITEX sets \sfcodes x=1000 for all x, except that \sfcodes X=999 for uppercase
% letters. The following changes are needed:

```

```
\sfcodes`\)=0 \sfcodes`\'=0 \sfcodes`\]=0
```

```
% The \nonfrenchspacing macro will make further changes to \sfcodes values.
```

Definitions like this depend on the language. Because original TeX was mostly meant for typesetting English, these things are hard coded. In ConTeXt such definitions relate to languages.

I show these definitions because they also illustrate what TeX is about: typesetting math:

```
% Finally, INITEX sets all \delcode values to -1, except \delcode`.=0
```

```
\delcode`\(\="028300
\delcode`\)=029301
\delcode`\[="05B302
\delcode`\]=05D303
\delcode`\<="26830A
\delcode`\>="26930B
\delcode`\/="02F30E
\delcode`\|="26A30C
\delcode`\\"="26E30F
```

```
% N.B. { and } should NOT get delcodes; otherwise parameter grouping fails!
```

Watch the last comment. One of the complications of TeX is that because some characters have special meanings, we also need to deal with exceptions. It also means that

arbitrary input is not possible. For instance, unless the percent character is made a letter, everything following it till the end of a line will be discarded. This is an area where macro packages can differ but in MkII we followed these rules. In MkIV we made what we called `\nonknuthmode` default which means that ampersands are just that and scripts are only special in math (there was also `\donknuthmode`). So, ConTeXt is not like plain there.

8.3 Housekeeping

The next section defines some numeric shortcuts. Here the fact is used that a defined symbolic character can act as counter value. When the number is larger than 255 a math character is to be used. In `Luatex`, which is a Unicode engine character codes can be much larger.

```
% To make the plain macros more efficient in time and space, several constant
% values are declared here as control sequences. If they were changed, anything
% could happen; so they are private symbols.
```

```
\chardef\@ne=1
\chardef\@tw@=2
\chardef\@thr@@=3
\chardef\@sixt@@n=16
\chardef\@cclv=255
\mathchardef\@cclvi=256
\mathchardef\@m=1000
\mathchardef\@M=10000
\mathchardef\@MM=20000
```

In ConTeXt we still support these shortcuts but never use them ourselves. We have plenty more variables and constants and nowadays always use verbose names. (There was indeed a time when each extra characters depleted string memory more and more so then using short command names made sense.) The comment is right that using such variables is more efficient, for instance once loaded a macro is a sequence of tokens, so `\@one` takes one memory slot. In the case of the first three the saving is zero and even interpreting a single character token `3` is not less efficient than `\@thr@@`, but in the case of `\@cclv` the three tokens `255` take more memory and also trigger the number scanner which is much slower than simply taking the meaning of the `\chardef`'d token. However, the ConTeXt variable `\plusone` is as efficient as the `\@one` and it looks prettier in code too (and I'm very sensitive for that). So, here ConTeXt is definitely different!

It makes no sense to show the next section here: it deals with managing registers, like counters and dimensions and token lists. Traditional `TeX` has 255 registers per category. Associating a control sequence (name) with a specific counter is done with `\countdef` but I don't think that you will find a macro package that expects a user to use that primitive. Instead it will provide a `\newcount` macro. So yes, here ConTeXt is like plain.

Understanding these macros is a test case for understanding \TeX . Take the following snippet:

```
\let\newtoks=\relax % we do this to allow plain.tex to be read in twice
\outer\def\newhelp#1#2{\newtoks#1#1\expandafter{\csname#2\endcsname}}
\outer\def\newtoks{\alloc@5\toks\toksdef@cclvi}
```

The `\outer` prefix flags macros as to be used at the outermost level and because the `\newtoks` is in the macro body of `\newtoks` it has to be relaxed first. Don't worry if you don't get it. In ConTeXt we have no outer macros so the definitions differ there.

The plain format assumes that the first 10 registers are used for scratch purposes, so best also assume this to be the case in other macro packages. There is no need for ConTeXt to differ from plain here. The definitions of box registers and inserts are special: there is no `\boxdef` and inserts use multiple registers. Especially the allocation of inserts is macro package specific. Anyway, ConTeXt users never see such details because inserts are used as building blocks deep down.

Right after defining the allocators some more constants are defined:

```
% Here are some examples of allocation.
```

```
\newdimen\maxdimen \maxdimen=16383.99999pt % the largest legal <dimen>
```

We do have that one, as it's again a standard but we do have more such constants. This definition is kind of interesting as it assumes knowledge about what is acceptable for \TeX as dimension:

```
{\dimen0=16383.99999pt \the\dimen0 \quad \number\dimen0}
{\dimen0=16383.99998pt \the\dimen0 \quad \number\dimen0}
```

```
16383.99998pt 1073741823
16383.99998pt 1073741823
```

Indeed it is the largest legal dimension but the real largest one is slightly less. We could also have said the following, which also indicates what the maximum cardinal is:

```
\newdimen\maxdimen \maxdimen=1073741823sp
```

We dropped some of the others defined in plain. So, ConTeXt is a bit like plain but differs substantially. In fact, MkII already used a different allocator implementation and MkIV is even more different. We also have more `\new` things.

The `\newif` definition also differs. Now that definition is quite special in plain \TeX , so if you want a challenge, look it up. It defines three macros as the comment says:

```
% For example, \newif\iffloor creates \foottrue, \footfalse to go with \iffloor.
```

The `\iffloor` is either equivalent to `\iftrue` or `\iffalse` because that is what \TeX needs to see in order to be able to skip nested conditional branches. In ConTeXt we have so

called conditionals, which are more efficient. So, yes, you will find such defined ifs in the ConTeXt source but way less than you'd expect in such a large macro package: ConTeXt code doesn't look much like plain code I fear.

8.4 Parameters

A next stage sets the internal parameters:

```
% All of TeX's numeric parameters are listed here, but the code is commented out
% if no special value needs to be set. INITEX makes all parameters zero except
% where noted.
```

We use different values for many of them. The reason is that the plain TeX format is set up for a 10 point Computer Modern font system, and for a particular kind of layout, so we use different values for:

```
\hsize=6.5in
\vsiz=8.9in
\maxdepth=4pt
```

and

```
\abovedisplayskip=12pt plus 3pt minus 9pt
\abovedisplayshortskip=0pt plus 3pt
\belowdisplayskip=12pt plus 3pt minus 9pt
\belowdisplayshortskip=7pt plus 3pt minus 4pt
```

No, here ConTeXt is not like plain. But, there is one aspect that we do inherit and that is the ratio. Here a 10 point relates to 12 point and this 1.2 factor is carried over in some defaults in ConTeXt. So, in the end we're a bit like plain.

After setting up the internal quantities plain does this:

```
\newskip\smallskipamount \smallskipamount=3pt plus 1pt minus 1pt
\newskip\medskipamount \medskipamount=6pt plus 2pt minus 2pt
\newskip\bigskipamount \bigskipamount=12pt plus 4pt minus 4pt
\newskip\normalbaselineskip \normalbaselineskip=12pt
\newskip\normallineskip \normallineskip=1pt
\newdimen\normallineskiplimit \normallineskiplimit=0pt
\newdimen\jot \jot=3pt
\newcount\interdisplaylinepenalty \interdisplaylinepenalty=100
\newcount\interfootnotelinepenalty \interfootnotelinepenalty=100
```

The first three as well as the following three related variables are not internal quantities but preallocated registers. These are not used in the engine but in macros. In ConTeXt we do provide them but the first three are never used that way. The last three are not defined at all. So, ConTeXt provides a bit what plain provides, just in case.

8.5 Fonts

The font section is quite interesting. I assume that one reason why some want to warn users against using ConTeXt is because it supports some of the font switching commands found in plain. We had no reasons to come up with different ones but they do different things anyway, for instance adapting to situations. So, in ConTeXt you will not find the plain definitions:

```
\font\tenrm=cmr10 % roman text
\font\preloaded=cmr9
\font\preloaded=cmr8
\font\sevenrm=cmr7
\font\preloaded=cmr6
\font\fiverm=cmr5
```

There is another thing going on here. Some fonts are defined `\preloaded`. So, `cmr9` is defined, and then `cmr8` and `cmr6`. But they all use the same name. Later on we see:

```
\let\preloaded=\undefined % preloaded fonts must be declared anew later.
```

If you never ran into the relevant part of the TeX book or read the program source of TeX, you won't realize that preloading means that it stays in memory which in turn means that when it gets (re)defined later, the font data doesn't come from disk. In fact, as the plain format is normally dumped for faster reload later on, the font data is also retained. So, preloading is a speed up hack. In ConTeXt font loading has always been delayed till the moment a font is really used. This permits plenty of definitions and gives less memory usage. Of course we do reuse fonts once loaded. All this, plus the fact that we have a system of related sizes, collections of families, support multiple font encodings alongside, collect definitions in so called typescript, etc. makes that the ConTeXt font subsystem is far from what plain provides. Only some of the command stick, like `\rm` and `\bf`.

The same is true for math fonts, where we can have different math font setups in one document. Definitely in MkII times, we also had to work around limitations in the number of available math families, which again complicated the code. In MkIV things are even more different, one can even consider the implementation somewhat alien for a standard macro package, but that's for another article (if at all).

8.6 Macros

Of course ConTeXt comes with macros, but these are organized in setups, environments, instances, etc. The whole process and setup is keyword driven. Out of the box all things work: nothing needs to be loaded. If you want it different, you change some settings, but you don't need to load something. Maybe that last aspect is what is meant with ConTeXt being like plain: you don't (normally) load extra stuff. You just adapt the system to your needs. So there we proudly follow up on plain TeX.

In the plain macro section we find definitions like:

```
\def\frenchspacing{\sfcode`\.@m \sfcode`?\@m \sfcode`!\@m
  \sfcode`:@\@m \sfcode`;\@m \sfcode`,,\@m}
\def\nonfrenchspacing{\sfcode`3000\sfcode`?3000\sfcode`!3000%
  \sfcode`:@2000\sfcode`@1500\sfcode`@1250 }
```

and:

```
\def\space{ }
\def\empty{ }
\def\null{\hbox{ }}

\let\bgroun{%
\let\egroun{}
```

and:

```
\def\nointerlineskip{\prevdepth-1000\p@}
\def\offinterlineskip{\baselineskip-1000\p@
 \lineskip\z@ \lineskip\lim\maxdimen}
```

Indeed we also provide these, but apart from the two grouping related aliases their implementation is different in ConTeXt. There is no need to reinvent names.

For a while we kept (and did in MkII) some of the plain helper macros, for instance those that deal with tabs, but we have several more extensive table models that are normally used. We always had our own code for float placement, and we also have more options there. Footnotes are supported but again we have multiple classes, placements, options, etc. Idem for itemized lists, one of the oldest mechanisms in ConTeXt. We don't have `\begin{section}` but of course we do have sectioning commands, and have no `\proclaim` but provide lots of descriptive alternatives, so many that I forgot about most of them by now (so plain is a winner in terms of knowing a macro package inside out).

The fact that we use tables, floats and footnotes indeed makes ConTeXt to act like plain, but that's then also true for other macro packages. A fact is that plain sets the standard for how to think about these matters! The same is true for naming characters:

```
\chardef\%="`%
\chardef\&="`&
\chardef\#=`\'#
\chardef\$="`$%
\chardef\ss="19%
\chardef\ae="1A%
\chardef\oe="1B%
\chardef\o="1C%
\chardef\AE="1D%
```

```
\chardef\OE="1E
\chardef\O="1F
\chardef\i="10 \chardef\j="11 % dotless letters
```

But we have many more and understandable the numbers are different in ConTeXt because we use different font (encodings). Their implementation is more adaptive. The same is true for accented characters:

```
\def`\#1{{\accent18 #1}}
\def'\#1{{\accent19 #1}}
```

The definitions in MkII are different (in most cases we use native glyphs) and in MkIV we use Unicode anyway. I think that the `\accent` command is only used in a few exceptional cases (like very limited fonts) in MkII and never in MkIV. The implementation of for instance accents (and other pasted together symbols) in math is also quite different.

There are also definitions that seem to be commonly used in macro packages but that we never use in ConTeXt because they interfere badly with all kind of other mechanisms, so you will find no usage of

```
\def\leavevmode{\unhbox\voidb@x} % begins a paragraph, if necessary
```

in ConTeXt. In order to stress that we provide `\dontleavehmode`, a wink to not using the one above.

The macro section ends with lots of math definitions. Most of the names used are kind of standard so again here ConTeXt is like plain, but the implementation can differ as does the level of control.

8.7 Output

Once a page is ready it gets wrapped up and shipped out. Here ConTeXt is very different from plain. The amount of code in plain is not that large but the possibilities aren't either, which is exactly what the objectives demand: a simple (example) format that can be described in the TeXbook. But, as with other aspects of plain, it steered the way macro packages started out as it showed the way. As did many examples in the TeX book.

8.8 Hyphenation

As an afterthought, the plain format ends with loading hyphenation patterns, that is the English ones. That said it will be clear that ConTeXt is not like plain: we support many languages, and the subsystem deals with labels, specific typesetting properties, etc. too.

```
\lefthyphenmin=2 \righthyphenmin=3 % disallow x- or -xx breaks
\input hyphen
```

We don't even use these patterns as we switched to utf long ago (also in MkII) if only because we had to deal with a mix of font encodings. But we did preload the lot there. In MkIV again things are much different.

8.9 Conclusion

The plain format does (and provides) what it is supposed to do. It is a showcase of possibilities and part of the specification. In that respect it's nice that ConTEXt is considered to be like plain. But if it wasn't more, there was no reason for its existence. Like more assumptions about ConTEXt it demonstrates that those coming up with answers and remarks like that probably missed something in assessing ConTEXt. Just let users find out themselves what suits best (and for some that actually might be plain \TeX).

Let me make one thing clear. If you look at the documents that describe the development of \TeX , MetaFont and the related fonts, you can only awe at what was done on hardware that doesn't come close to what we hold now in the palm of our hand. And it was done in a relative short time span. The fact that plain \TeX ran on it the way it did is amazing. Anyone who considers criticizing \TeX and plain \TeX should think (more than) twice.

9.1 Introduction

Hans Hagen
Hasselt NL
July 2019 (public 2023)

The \TeX typesetting system is pretty powerful, and even more so when you combine it with MetaPost and Lua. Add an xml parser, a whole lot of handy macros, provide support for fonts and advanced pdf output and you have a hard to beat tool. We're talking Con \TeX xt.

Such a system is very well suited for fully automated typesetting. There are \TeX lovers who claim that \TeX can do anything better than the competition but that's not true. Automated typesetting is quite doable when you accept the constraints. When the input is unpredictable you need to play safe!

Some things are easy: turning complex xml into pdf with adaptive graphics, fast data processing, colorful layouts, conditional processing, extensive cross referencing, you can safely say that it can be done. But in practice there is some design involved and those are often specified by people who manipulate a layout on the fly and tweak and cheat in an interactive wysiwyg program. That is however not an option in automated typesetting. Traditional thinking with manual intervention has to make place for systematic and consistent solutions. Limitations can be compensated by clever designs and getting the maximum out of the system used.

Unfortunately in practice some habits are hard to get rid of. Inconsistent use of colors, fonts, sectioning, image placements are just a few aspects that come to mind. When you typeset educational documents you also have to deal with strong opinions about how something should be presented and what students can't (!) handle, like for instance cross references. One of the most dominant demands in typesetting such documents are so called side floats. In (for instance) scientific publishing references to content typeset elsewhere (formulas, graphics) is acceptable but in educational documents this is often not an option (don't ask me why).

In the next sections I will mention a few aspects of side floats. I will not discuss the options because these are covered in manuals. Here we stick to the challenges and the main question that you have to ask yourself is: "How would I solve that if it can be solved at all?". It might make you a bit more tolerant for suboptimal outcome.

9.2 The basics

We start with a simple example. The result is shown in figure 9.1. We have figures, put at the left, with enough text alongside so that we don't have a problem running into the next figure.

```
\dorecurse {8} {
    \useMPlibrary[dum]
    \setuplayout[middle]
    \setupbodyfont[plex]
    \startplacefigure[location=left]
        \externalfigure[dummy] [width=3cm]
    \stopplacefigure
    \samplefile{sapolsky}
    \par
}
```

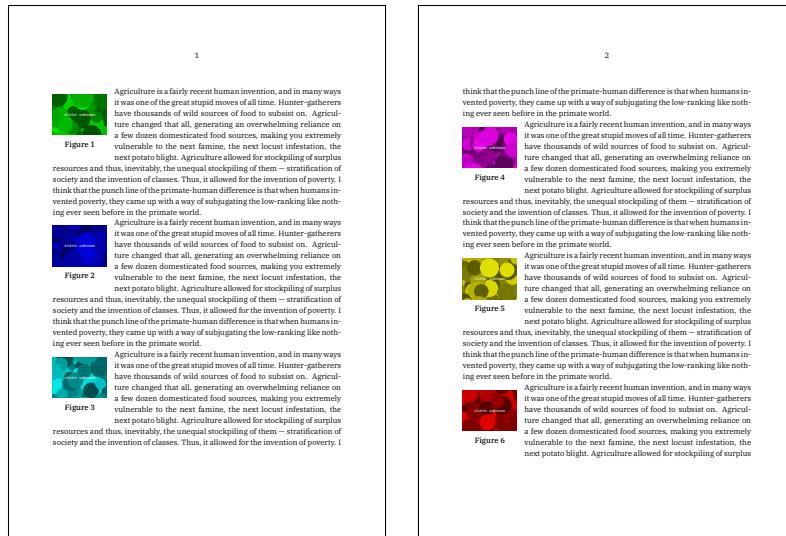


Figure 9.1 A simple example with enough text in a single paragraph.

Challenge: Anchor some boxed material to the running text and make sure that the text runs around that material. When there is not enough room available on the page, enforce a page break and move the lot to the next page.

But more often than not, the following paragraph is not long enough to go around the insert. The worst case is of course when we end up with one word below the insert, for which the solution is to adapt the text or make the insert wider or narrower. Forgetting about this for now, we move to the case where there is not enough text: figure 9.2.

```
\dorecurse {8} {
    \useMPlibrary[dum]
    \setuplayout[middle]
    \setupbodyfont[plex]
    \startplacefigure[location=left]
```

```

\externalfigure[dummy] [width=3cm]
\stopplacefigure
\samplefile{ward} \par \samplefile{ward}
\par
}

```

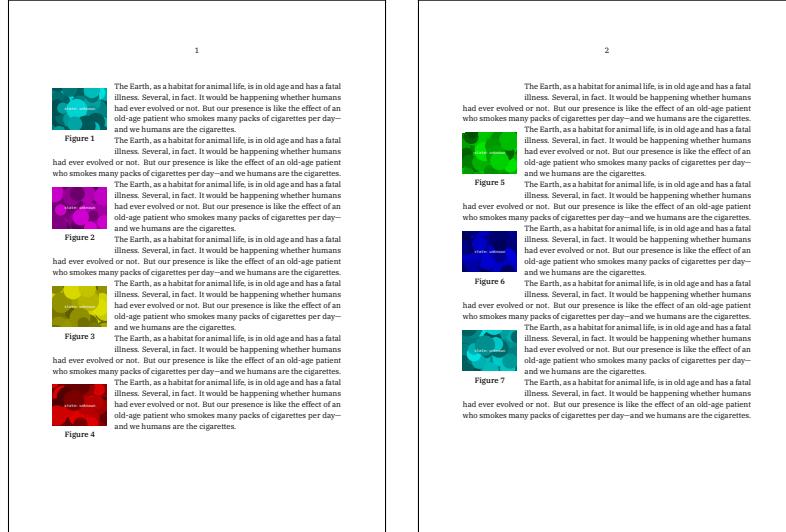


Figure 9.2 A simple example with enough text but multiple paragraphs.

Challenge: At every new paragraph, check if we're still not done with the blob we're typesetting around and carry on till we are behind the insert.

The next example, shown in figure 9.3, has less text. However, the running text is still alongside the figure, so this means that white space need to be added till we're beyond.

```

\dorecurse {8} {
\useMPlibrary[dum]
\setuplayout[middle]
\setupbodyfont[plex]
\startplacefigure[location=left]
\externalfigure[dummy] [width=3cm]
\stopplacefigure
\samplefile{ward}
\par
}

```

Challenge: When there is not enough content, and the next insert is coming, we add enough whitespace to go around the insert and then start the new one. This is typically something that can also be enforced by an option.

Before we move on to the next challenge, let's explain how we run around the insert. When TeX typesets a paragraph, it uses dimensions like `\leftskip` and `\rightskip` (margins) and shape directives like `\hangindent` and `\hangafter`. There is also the possibility to define a `\parshape` but we will leave that for now. The width of the image

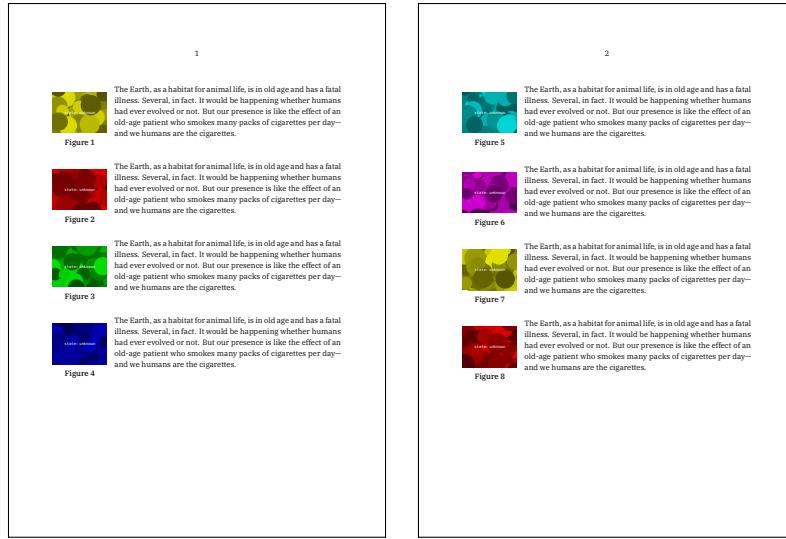


Figure 9.3 A simple example with less text

is reflected in the indent and the height gets divided by the line height and becomes the `\hangafter`. Whenever a new paragraph is started, these parameters have to be set again.⁵ In ConTeXt hanging is also available as basic feature.

```
\starthanging[location=left]
  {\blackrule[color=maincolor, width=3cm, height=1cm]}
  \samplefile{carrol}
\stophanging
```

The fraction of fossil olfactory receptor genes is significantly higher in all species with full color vision. This suggests that the evolution of trichromatic vision — which allows these primates to detect food, mates, and danger with visual cues — has reduced their reliance on the sense of smell.

```
\starthanging[location=right]
  {\blackrule[color=maincolor, width=10cm, height=1cm]}
  \samplefile{jojomayer}
\stophanging
```

If we surrender the thing that separates us from machines, we will be replaced by machines. The more advanced machines will be, the more human we will have to become.

The hanging floats are not implemented this way but are hooked into the paragraph start routines. The original approach was a variant of the macros by Daniel Comenetz as published in TUGBoat Volume 14 (1993), No. 1: Anchored Figures at Either Margin. In the meantime they are far from that, so ConTeXt users can safely blame me for any issues.

⁵ I still consider playing with a third parameter representing hang height and add that to the line break routine, but I have to admit that tweaking that is tricky. Do I really understand what is going on there?

9.3 Unpredictable dimensions

In an ideal world images will be sort of consistent but in practice the dimension will differ, even fonts used in graphics can be different, and they can have white space around them. When testing a layout it helps to use mockups with a clear border. If these look okay, one can argue that worse looking assemblies (more visual whitespace above of below) is a matter of making better images. In figure 9.4 we demonstrate how different dimensions influence the space below the placement.

```
\dostepwiserecurse {2} {8} {1} {
  \useMPlibrary [dum]
  \setuplayout [middle]
  \setupbodyfont [plex]
  \setupalign [tolerant,stretch]
  \startplacefigure [location=left]
    \externalfigure [dummy] [width=#1cm]
  \stopplacefigure
  \samplefile{sapolsky}
  \par
}
```

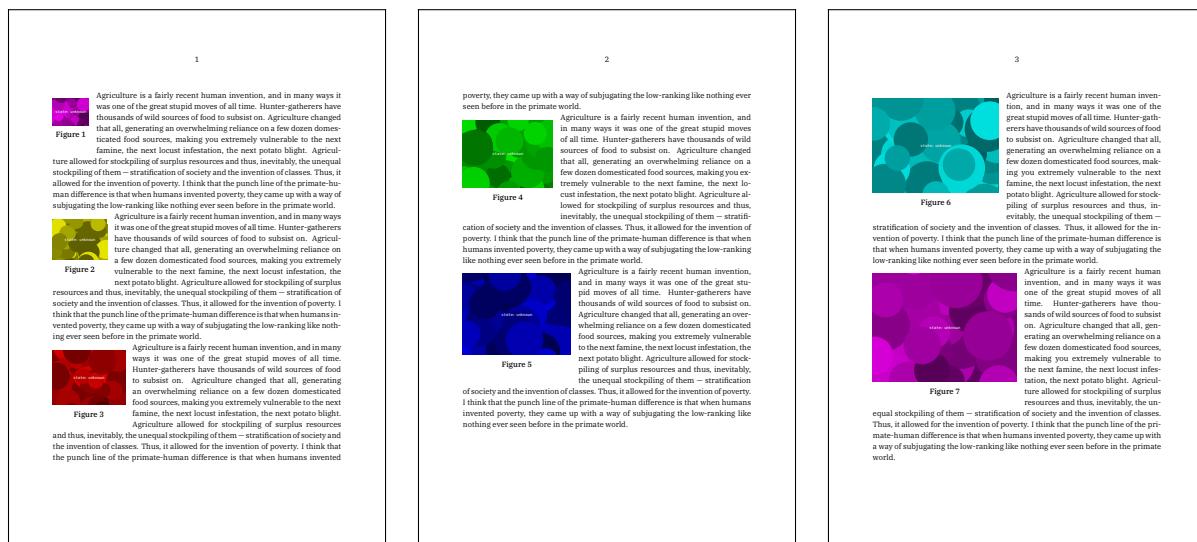


Figure 9.4 Spacing relates to dimensions.

In ConTeXt there are plenty of options to add more space above or below the image. You can anchor the image to the first line in different ways and you can move it some lines down, either or not with text flowing around it. But here we stick to simple cases, we only discuss the challenges.

Challenge: Adapt the wrapping to the right dimensions and make sure that the (optional) caption doesn't overlap with the text below.

9.4 Moving forward

When the insert doesn't fit it has to move, which is why it's called a float. One solution is to take it out of the page stream and turn it into a regular placement, normally centered horizontally somewhere on the page, and in this case probably at the top of one of the next pages. Because we can cross reference this is a quite okay solution. But, in educational documents, where authors refer to the graphic (picture) on the left or right, that doesn't work out well. The following content is bound to the image.

Calculating the amount of available space is a bit tricky due to the way \TeX works. But let's assume that this can be done, in Con \TeX t we have seen several strategies for this, we then end up at the top of the next page and there different spacing rules apply, like: no spacing at the top at all. In our examples no whitespace between paragraphs is present. The final solutions are complicated by the fact that we need to take this into account.

Challenge: Make sure that we never run off the page but also that we don't end up with weird situations at the top of the next page.

Another possibility is that images so tightly fit a whole number of lines, that a next one can come too close to a previous one. Again, this demands some analysis. Here we use examples with captions but when there are no captions, there is also less visual space (no depth in lines).

Challenge: Make sure that a following insert never runs too close to a previous insert.

Solutions can be made better when we use multi-pass information. Because in a typical \TeX run there is only looking back, storing information can actually make us look forward. But, as in science fiction: when you act upon the future, the past becomes different and therefore also the future (after that particular future). This means that you can only go forward. Say you have 10 cases: when case 5 changes because of some feedback, then case 6 upto 10 also can change. So, you might need more than 10 runs to get things right. In a workflow where users are waiting for a result, and a few hundred side floats are used this doesn't sell well: processing 400 pages with a 20 page per second rate takes 20 seconds per run. Normally one needs a few runs to get the references right. Assuming a worst case of 60 seconds, 10 extra runs will bring you close to 15 minutes. No deal.

Of course one can argue for some load-in-memory and optimize in one go, but although \TeX can do that well for paragraphs, it won't work for complex documents. Sure, it's a nice academic exercise to explore limited cases but those are not what we encounter.

9.5 Cooperation

When discussing (on YouTube) "Extending Darwin's Revolution" David Sloan Wilson and Robert Sapolsky touch on the fact that in some disciplines (like economics) evolutionary principles are applied. One can apply for instance the concept of a 'selfish

gene'. However, they argue that when doing that, one actually lags behind the now accepted group selection (which goes beyond the individual benefits). An example is given where aggressive behavior on the short term can turn one in a winner (who takes it all) but which can lead to self destructive in the long run: cooperating seems to work better than terminal competition.

In \TeX we have glues and penalties. The machinery likes to break at a glue but a severe penalty can prohibit that. The fact that we have penalties and no rewards is interesting: a break can be stimulated by a negative penalty. I've forgotten most of what I learned about cognitive psychology but I do remember that penalty vs reward discussions could get somewhat out of hand.

So, when we have in the node list a mix of glue (you can break here), penalties (better not break here) and rewards (consider breaking here) you can imagine that these nodes compete. The optimal solution is not really a group process but basically a rather selfish game. Building a system around that kind of cooperation is not easy. In ConTeXt a lot of attention always went into consistent vertical spacing. In MkII there were some 'look back' and 'control forward' mechanisms in place, and in MkIV we use a model of weighted glue: a combination of penalties and skips. Again we look back and again we also try to control the future. This works reasonable well but what if we end up in a real competition?

A section head should not end up at the bottom of a page. Because when it gets typeset it is unknown what follows, it does some checking and then tries to make sure that there is no page break following. Of course there needs to be a provision for the cases that there are many (sub)heads and of course when there are only heads on a page (in a concept for instance) you don't want to run of the page.

Similar situations arise with for instance itemized lists and the tabulate mechanism. There we have some heuristics that keep content together in a way that makes sense given the construct: no single table line at the bottom of a page etc. But then comes the side float. The available space is checked. When doing that the whitespace following the section head has to collapse with the space around the image, but of course at the top of a page spacing is different. So, calculations are done, but even a small difference between what is possible and what is needed can eventually still trigger an unwanted page break. This is because you cannot really ask how much has been accumulated so far: the space used is influenced by what comes next (like whitespace, maybe interline space, the previous depth correction, etc). That in turn means that you have to (sort of) trigger these future space related items to be applied already.

Challenge: Let the side float mechanism nicely cooperate with other mechanisms that have their own preferences for crossing pages, adding whitespace and being bound to following content.

9.6 Easy bits

Of course, once there is such a mechanism in place, user demands will trigger more features. Most of these are actually not that hard to deal with: renumbering due to

moved content, automatic anchoring to the inner or outer margin, horizontal placement and shifting into margins, etc. Everything that doesn't relate to vertical placement is rather trivial to deal with, especially when the whole infrastructure for that is already present (as in ConTeXt). The problem with such extensions is that one can easily forget what is possible because most are rarely used.

Challenge: Make sure that all fits into an understandable model and is easy to control.

9.7 Conclusion

The side float mechanism in ConTeXt is complex, has many low level options, and its code doesn't look pretty. It is probably the mechanism that has been overhauled and touched most in the code base. It is also the mechanism that (still) can behave in ways you don't expect when combined with other mechanisms. The way we deal with this (if needed) is to add directives to (in our case) xml files that tells the engine what to do. Because that is a last resort it is only needed when making the final product. So in the end, we're still have the benefits of automated typesetting.

Of course we can come up with a different model (basically re-implement the page builder) but apart from much else falling apart, it will just introduce other constraints and side effects. Thinking in terms of selfish nodes, glues and penalties, works ok for a specific document where one can also impose usage rules. If you know that a section head is always followed by regular text, things become easier. But in a system like ConTeXt you need to update your thinking to group selection: mechanisms have to work together and that can be pretty complicated. Some mechanisms can do that better than others. One outcome can be that for instance side floats are not really group players, so eventually they might become less popular and fade away. Of course, as often, years later they get rediscovered and the cycle starts again. Maybe a strong argument can be made that in fully automated typesetting concepts like side floats should not be used anyway.

If I have to summarize this wrap up, the conclusion is that we should be realistic: we're not dealing with an expert system, but with a bunch of heuristics. You need an intelligent system to help you out of deadlock and oscillating solutions. Given the different preferences you need a multiple personality system. You might actually need a system that wraps your expectations and solutions and that adapts to changes in those over time. But if there is such a system (some day) it probably doesn't need you. In fact, maybe even typesetting is not needed any more by then.

10.1 Introduction

I'm always puzzled when I read that someone wonders if ConTeXt is still up to date or maintained because some manual has a timestamp of a decade ago. I'm also puzzled by some rants you can run into when searching the web. In the next few paragraphs I'll comment on this.

10.2 Stability

Say that you're an enthusiastic user of console commands like `ls` (`dir`), `cp` (`copy`) or maybe `ssh`, `rsync`, `curl`. How often do you consult a manual on how they evolve? And say that you, for some reason, do consult a manual, there is a good chance that it is pretty old. Does that mean that the commands are obsolete? The binaries probably get fixed for bugs but the interface stays the same, which is what you expect. Every time we generate a zip for the ConTeXt distribution, the related website also gets generated, using a bunch of xml files that get transformed to html using xslt and a pretty ancient version of `xsltproc` (why should I update). I never check for a new manual as it keeps doing the job. And additional manuals and reports get added.

So, once some functionality is stable, and a lot of macro code in ConTeXt is just that, there is no need to update a manual! Putting a new time stamp on it is basically fake updating. And often the more introductory kind of manuals don't need to be updated at all, apart from maybe cultural changes that demand a (political correct) update. Them being a bit old and not being updated is actually a good thing as it signals stability.

It is worth mentioning that the ConTeXt distribution is not the only source of information. There are manuals written by others and there is the Wiki. All is the work of volunteers and updating all that depends on how much time one can allocate.

10.3 Excuses

It is a fact that ConTeXt evolves. New functionality gets added and some mechanism get extended. Often these are described in dedicated manuals or articles that end up in collections, and there are plenty of them in the distribution. For some reason those complaining about a beginners manual with an old time stamp don't check if there is more, and there is quite some more! Don't only look at the ConTeXt garden (the wiki) but also keep an eye on what gets distributed. Some users are very good in track of what gets added, because sometimes I get fixes for typos send within a few hours after uploading a zip.

We appreciate that other users point out that writing manuals takes time and that indeed our time is not without limits. If I could sit down and write manuals whole day, and it would get paid, I might do it. But it is a fact that development of ConTeXt is not paid for at all. I can work on it in company time but much happens in spare time. Most development is a gamble on future use or done because we want to be complete or because code can be improved. So, writing a manual then closely relates to what we like doing: it determines the topics and priorities. If something gets explored and ends up in new functionality then that gets documented in the process. It is the fun factor that drives it. The same is true for LuaTeX development.

So, we have as valid excuse that new manuals relate to (new) functionality and old ones stay as they are. Don Knuth remarks somewhere that writing a manual as part of the development is a good thing. We fully agree with that.

10.4 Cutting edge

Does an old manual indicate that nothing happens? Definitely not. Over a decade of LuaTeX development is closely related to ConTeXt and there is plenty of reporting about that. Does that mean that we need to rewrite manuals? No, existing functionality remains. And of course users are free to come up with more detailed manuals (which they seldom do). Some developments get published in user group journals but we don't publish much about specific ConTeXt features and usage because it's hard to do that for a diverse audience.

Currently we have what is called ConTeXt MkXL (aka LMTX), but we also have the prelude to that, MkIV, and the frozen predecessor MkII. Apart from changes in technology (most noticeably fonts and encodings) the functionality is accumulative: most old manuals (unless they are specialized into old school fonts for instance) apply to the latest greatest version.

It is a misunderstanding that the development of ConTeXt, LuaTeX and LuaMetaTeX is somehow funded by projects that we do. This is not true. We can apply both in projects but as we charge by the hour (or day) no customer ever sees development on the bill. Of course during a project we can gain on efficiency (so then development pays back) and because we know the system style writing is efficient too. In fact, in most cases our customers don't know or care what tool we use because tools are expected to be part of the deal. Most projects we can (and could) only do because we can use ConTeXt and that is a side effect of the fact that we do develop beforehand. We're often the only technically and/or affordable way out. It's a chicken-egg issue: we have a tool and therefore get a project. We never get a projects where we can develop a tool. No one pays for TeX development or at least no one ever came to us with specific TeX related demands. It looks like the world takes it for granted that TeX is just there.⁶

⁶ There are a few subcomponents of ConTeXt that were partially sponsored by users and we do have some support contracts that permit experiments and development.

The reason for ConTeXt being cutting edge (in terms of TeX) is that we like challenges, that users demand features that are interesting to explore and that we've been part of the TeX scenery for a while now. We just like that.

It's good to know that ConTeXt was and is developed as a toolkit. We started long ago because we needed a way to quickly create and update reports of meetings that we chaired. Next we needed a way to efficiently produce high quality education materials (of various kind) and support maintenance of sets of related (quality assurance) manuals. We could have used wordstar, wordperfect or msword but liked the TeX way much more. As said, most customers didn't even know or care what tool was used because the (often highly interactive pdf) outcome mattered most. In fact, we would not be interested in this kind of work if we were forced to use clumsy tools, but for sure a lot can be done with those as well.

10.5 Continuity

Most development happens at Pragma ADE by me (Hans) with help from my two colleagues (Ton and Kees) and the community (Aditya, Alan, Mikael, Mojca, Luigi, Hraban, Taco, Thomas, Tomas, Willi, Wolfgang, and others).⁷ I won't mention those on the mailing list who contribute with ideas, testing and support, but they can't be missed. The biggest danger for continuity is a polluted code base where everyone just pushed code into a repository. So this is closely guarded. A user patch might work well for that user but can break something else.

With TeX you need to keep in mind that once a solution works there is no need to update code or manuals. As long as there is a working LuaTeX (LuaMetaTeX) binary you're fine. Maybe if some specific fonts are used, a filename might need to be adapted.

An example. When we added a new xml subsystem to ConTeXt MkIV we knew that some day we could use it. We now use it in a few projects and I'm pretty sure that we would not do these projects otherwise as it would demand writing quite complex xslt style sheets that then would have to be applied to thousands of files per run. To some extend what is available in ConTeXt sort of drives the kind of work you look for. That said: if you consider using ConTeXt for simple or complex documents, either or not in a collection, either or not using TeX or xml input you can be assured that this will work (and might even get better) because we use it ourselves.

If you want to get an idea about development, just look at the (five) documents that describe the development of LuaTeX (LuaMetaTeX). Locating them in the distribution is a good opportunity to explore the documents. They will show you what happened the last decade(s) and give you some trust in ConTeXt. Or come to a ConTeXt meeting and meet those involved.

⁷ More names could be here as I write this in 2022.

10.6 Closing remark

So next time someone asks if ConTeXt is maintained because some old manual stays around, return the question if frequently updated manuals are a sign of stability. Also ask if someone looked a bit in the documentation tree. The oldest manual in the TeX world is the TeXbook that describes the oldest stable set of macros: plain TeX. There are happy users out there who love that stability. If it were not for the wonderful personality of Don Knuth this program would already been forgotten. I think that long term stability and unchanged code and manuals are something that we need to cherish and get accustomed to, which is not easy in a time when a phone and its operating system are outdated as soon as you unbox it. It's also not easy in a time of instant communication, more and more confused by what is called artificial intelligent mumbling, but that's for another wrapup.

Hans Hagen
Hasselt NL

(uncorrected so there's something left to complain)

Performance again 11

Hans Hagen
Hasselt NL
Februari 2020 (public 2023)

11.1 Introduction

In a *Maps* article of 2019 I tried to answer the question ‘Is \TeX really slow?’. A while after it was published on the Dutch \TeX mailing list a user posted a comment stating that in his experience the LuaTeX engine in combination with \LaTeX was terribly slow: one page per second for a Japanese text. It was also slower than pdfTeX with English, but for Japanese it was close to unusable. The alternative, using a Japanese \TeX engine was no option due to lack of support for certain images.

In order to check this claim I ran a test in ConTeXt . Even on my 8 year old laptop I could get 45 pages per second for full page Japanese texts (6 paragraphs with each 300 characters per page): 167 pages took just less than 4 seconds. Typesetting Japanese involves specific spacing and line break handling. So, naturally the question arises: why the difference. Frans Goddijn wondered if I could explain a bit more about that, so here we go.

In the mentioned article I already have explained what factors play a role and the macro package is one of them. It is hard to say to what extent inefficient macros or a complex layout influence the runtime, but my experience is that it is pretty hard to get speeds as low as 1 page per second. On an average complex document like the LuaTeX manual (lots of verbatim and tables, but nothing else demanding apart from color being used and a unique MetaPost graphic per page) I get at least a comfortable 20 pages per second.

I can imagine that for a \TeX user who sees other programs on a computer do complex things fast, the performance of \TeX is puzzling. But, where for instance rendering videos can benefit from specific features of (video) processors, multiple cores, or just aggressive optimization by compilers of (nested) loops and manipulation of arrays of bytes, this is not the case for \TeX . This program processes all in sequence, there is not much repetition that can be optimized, it cannot exploit the processor in special ways and the compiler can not do that many optimizations.

I can’t answer why a \LaTeX run is slower than a ConTeXt run. Actually, one persistent story has always been that ConTeXt was slow in comparison. But maybe it helps to know a bit what happens deep down in \TeX and how macro code can play a role in performance. When doing that I will simplify things a bit.

11.2 Text and nodes

The \TeX machinery takes input and turns that into some representation that can be turned into a visual representation ending up as pdf. So say that we have this:

`hello`

In a regular programming language this is a string with five characters. When the string is manipulated it is basically still a sequence of bytes in memory. In \TeX , if this is meant as text, at some point the internal representation is a so called node list:

`[h] -> [e] -> [l] -> [l] -> [o]`

In traditional \TeX these are actually character nodes. They have a few properties, like what font the character is from and what the character code is (0 up to 255). At some point \TeX will turn that list into a glyph list. Say that we have this:

`efficien`

This will eventually become seven nodes:

`[e] -> [ffi] -> [c] -> [i] -> [e] -> [n] -> [t]`

The ffi ligature is a glyph node which actually also keeps information about this one character being made from three.

In LuaTeX it is different, and this is one of the reasons for it being slower. We stick to the first example:

`[h] <-> [e] <-> [l] <-> [l] <-> [o]`

So, instead of pointing to the next node, we also point back to the previous: we have a double linked list. This means that all over the program we need to maintain these extra links too. They are not used by \TeX itself, but handy at the Lua end. But, instead of only having the font as property there is much more. The \TeX program can deal with multiple languages at the same time and this relates to hyphenation. In traditional \TeX there are language nodes that indicate a switch to another language. But in LuaTeX that property is kept with each glyph node. Actually, even specific language properties like the hyphen min, hyphen max and the choice if uppercase should be hyphenated are kept with these nodes. Spaces are turned into glue nodes, and these nodes are also larger than in regular \TeX engines.

So, in LuaTeX , when a character goes from the input into a node, a more complex data structure has to be set up and the larger data structure also takes more memory. That in turn means that caching (close to the cpu) gets influenced. Add to that the fact that we operate on 32 bit character values, which also comes with higher memory demands.

We mentioned that a traditional engine goes from one state of node list into another (the ligature building). Actually this is an integrated process: a lot happens on the

fly. If something is put into a `\hbox` no hyphenation takes place, only ligature building and kerning. When a paragraph is typeset, hyphenation happens on demand, in places where it makes sense.

In `LuaTeX` these stages are split. A node list is *always* hyphenated. This step as well as ligature building and kerning are *three* separate steps. So, there's always more hyphenation going on than in a traditional `TeX` engine: we get more discretionary nodes and again these take more memory than before; also the more nodes we have, the more it will impact performance down the line. The reason for this is that each step can be intercepted and replaced by a Lua driven one. In practice, with modern OpenType fonts that is what happens: these are dealt with (or at least managed in) Lua. For Japanese for sure the built-in ligature and kerning doesn't apply: the work is delegated and this comes at a price. Japanese needs no hyphenation but instead characters are treated with respect to their neighbors and glue nodes are injected when needed. This is something that Lua code is used for so here performance is determined by how well the plugged in code behaves. It can be inefficient but it can also be so clever that it just takes a bit of time to complete.

I didn't mention another property of nodes: attributes. Each node that has some meaning in the node list (glyphs, kerns, glue, penalties, discretionary, . . . , these terms should ring bells for a `TeX` user) have a pointer to an attribute list. Often these are the same for neighboring nodes, but they can be different. If a macro package sets 10 attributes, then there will be lists of ten attributes nodes (plus some overhead) active. When values change, copies are made with the change applied. Grouping even complicates this a little more. This has an impact on performance. Not only need these lists be managed, when they are consulted at the Lua end (as they are meant as communication with that bit of the engine) these lists are interpreted. It all adds up to more runtime. There is nothing like that in traditional `TeX`, but there some more macro juggling to achieve the same effects can cause a performance hit.

11.3 Macros and tokens

When you define a macro like this:

```
\def\MyMacro#1{\hbox{here: #1!}}
```

the `TeX` engine will parse this as follows (we keep it simple):

```
\def      primitive token
\MyMacro user macro pointing to:
  #1      argument list of length 1 and no delimiters
  {      openbrace token
  \hbox  hbox primitive token
    h    letter token h
    e    letter token e
    r    letter token r
```

```

e      letter token e
:
:      other token :
space token
#1     reference to argument
!
!      other token !
}      close brace token

```

The `\def` is eventually lost, and the meaning of the macro is stored as a linked list of tokens that get bound to the user macro `\MyMacro`. The details about how this list is stored internally can differ a bit per engine but the idea remains. If you compare tokens of a traditional \TeX engine with \LaTeX , the main difference is in the size: those in \LaTeX take more memory and again that impacts performance.

11.4 Processing

Now, for a moment we step aside and look at a regular programming language, like Pascal, the language \TeX is written in, or C that is used for \LaTeX . The high level definitions, using the syntax of the language, gets compiled into low level machine code: a sequence of instructions for the cpu. When doing so the compiler can try to optimize the code. When the program is executed all the cpu has to do is fetch the instructions, and execute them, which in turn can lead to fetching data from memory. Successive versions of cpu's have become more clever in handling this, predicting what might happen, (pre) fetching data from memory etc.

When you look at scripting languages, again a high level syntax is used but after interpretation it becomes compact so called bytecode: a sequence of instructions for a virtual machine that itself is a compiled program. The virtual machine fetches the bytes and acts upon them. It also deals with managing memory and variables. There is not much optimization going on there, certainly not when the language permits dynamically changing function calls and such. Here performance is not only influenced by the virtual machine but also by the quality of the original code (the scripts). In \LaTeX we're talking Lua here, a scripting language that is actually considered to be pretty fast.

Sometimes bytecode can be compiled Just In Time into low level machine code but for \LaTeX that doesn't work out well. Much Lua code is executed only once or a few times so it simply doesn't pay off. Apart from that there are other limitations with this (in itself impressive) technology so I will not go into more detail.

So how does \TeX work? It is important to realize that we have a mix of input and macros. The engine interprets that on the fly. A character enters the input and \TeX has to look at it in the perspective of what it expects. Is it just a character? Is it part of a control sequence that started (normally) with a backslash? Does it have a special meaning, like triggering math mode? When a macro is defined, it gets stored as a linked list of tokens and when it gets called the engine has to expand that meaning. In the process some actions themselves kind of generate input. When that happens a new level of input is entered and further expansion takes place. Sometimes \TeX looks ahead and

when not satisfied, pushes something back into the input which again introduces a new level. A lot can happen when a macro gets expanded. If you want to see this, just add `\tracingall` at the top of your file: you will be surprised! You will not see how often tokens get pushed and popped but you can see how much got expanded and how often local changes get restored. By the way, here is something to think about:

```
\count4=123
\advance \count4 by 123
```

If this is in your running text, the scanner sees `\count` and then triggers the code that handles it. That code expects a register number, here that is the `4`. Then it checks if there is an optional `=` which means that it has to look ahead. In the second line it checks for the optional keyword `by`. This optional scanning has a side effect: when the next token is *not* an equal or keyword, it has to push back what it just read (we enter a new input level) and go forward. It then scans a number. That number ends with a space or `\relax` or something not being a number. Again, some push back onto the input can happen. In fact, say that instead of `4` we have a macro indicating the register number, intermediate expansion takes place. So, even these simple lines already involve a lot of action! Now, say that we have this

```
% \newcounter \scratchcounter % done once
\scratchcounter 123
\scratchcounter =123
\advance\scratchcounter by 123
\advance\scratchcounter 123
```

Can you predict what is more efficient? If this operation doesn't happen frequently, performance wise there is no real difference between the variants with and without `=` and with and without `by`. This is because `\TeX` is pretty fast in tokenizing its input and interpreting its already stored token lists that have these commands. But given what we said before, when you talk of millions of such assignments, adding the equal sign and `by` could actually be faster because there is no pushing back onto the input stack involved. It probably makes no sense to take this into account when writing macros but just keep in mind that performance is in the details.

Actually, contrary to what you might expect, `\scratchcounter` is not even a counter in `\TeX`, and in `\LuaTeX` we can also do this:

```
% \newinteger\scratchcounter % done once
\scratchcounter 123
\scratchcounter =123
\advanceby\scratchcounter 123
```

Which means that because this counter is defined as so called “constant integer” it avoids some indirectness (to a counter register) and because `\advanceby` doesn't scan for a keyword the code above runs faster anyway.

This model of expansion is very different from compiled code or bytecode. To some extent you can consider a list of tokens that make up a macro to be bytecode, but instead of a sequence of bytes it is a linked list. That itself has a penalty in performance. Depending on how macros expand, the engine can be hopping all over the token memory following that list. That means that quite likely the data that gets accessed is not in your cpu cache and as a result performance cannot benefit from it apart of course from the expanding machinery itself, but that one is not a simple loop messing around with variables: it accesses code all over the place! Text gets hyphenated, fonts get applied, material gets boxed, paragraphs constructed, pages built. We're not moving a blob of bits around (as in a video) but we're constantly manipulating small amounts of memory scattered around memory space.

Now, where a traditional \TeX engine works on 8 bit characters and smaller tokens, the 32 bit \LaTeX works on larger chunks. Although macro names are stored as single symbolic units, there are moments when its real (serialized to characters) name is used, for instance when with `\csname`. When that happens, the singular token becomes a list, so for instance the (stored) token `\foo` becomes a temporary three token list (actually four if you also count the initial reference token). Those three tokens become three characters in a string that then is used in the hash lookup. There are plenty cases where such temporary string variables are allocated and filled. Compare:

```
\def\foo{\hello}
```

Here the macro `\foo` has just a one token reference to `\hello` because that's how a macro reference gets stored. But in

```
\def\foo{\csname hello\endcsname}
```

we have two plus five tokens to access what effectively is `\hello`. Each character token has to be converted to a byte into the assembled string. Now it must be said that in practice this is still pretty fast but when we have longer names and especially when we have utf8 characters in there it can come at a price. It really depends on how your macro package works and sometimes you just pay the price of progress. Buying a faster machine is then the solution because often we're not talking of extreme performance loss here. And modern cpu's can juggle bytes quite efficiently. Actually, when we go to 64 bit architectures, \LaTeX 's data structures fit quite well to that. As a side note: when you run a 32 bit binary on a 64 bit architecture there can even be a price being paid for that when you use \LaTeX . Just move on!

11.5 Management

Before we can even reach the point that some content becomes typeset, much can happen: the engine has to start up. It is quite common that a macro package uses a memory dump so that macros are not to be parsed each run. In traditional engines hyphenation patterns are stored in the memory dump as well. And some macro packages can put fonts in it. All kind of details, like upper- and lowercase codes can get stored too. In

LuaTeX fonts and patterns are normally kept out of the dump. That dump itself is much larger already because we have 32 bit characters instead of 8 bit so more memory is used. There are also new concepts, like catcode tables that take space. Math is 32 bit too, so more codes related to math are stored. Actually the format is so much larger that LuaTeX compresses it. Anyway, it has an impact on startup time. It is not that much, but when you measure differences on a one page document the overhead in getting LuaTeX up and running will definitely impact the measurement.

The same is true for the backend. A traditional engine uses (normally) Type1 fonts and LuaTeX relies on OpenType. So, the backend has to do more work. The impact is normally only visible when the document is finalized. There can be a slightly larger hickup after the last page. So, when you measure one page performance, it again pollutes the page per second performance.

11.6 Summary

So, to come back to the observation that LuaTeX is slower than pdfTeX . At least for ConTeXt we can safely conclude that indeed pdfTeX is faster when we talk about a standard English document, with TeX ascii input, where we can do with traditional small fonts, with only some kerning and simple ligatures. But as soon as we deal with for instance xml, have different languages and scripts, have more demanding layouts, use color and images, and maybe even features that we were not aware of and therefore didn't require in former times the LuaTeX engine (and for ConTeXt it's LuaMetaTeX follow up) performs way better than pdfTeX . And how about support for hyper links, protrusion and expansion, tagging for the sake of accessibility, new layers of abstraction, etc. The impact on performance can differ a lot per engine (and probably also per macro package). So, there is no simple answer and explanation for the fact that the observed slow LATEX run on Japanese text, apart from that we can say: look at the whole picture: we have more complex tokens, nodes, scripts and languages, fonts, macros, demands on the machinery, etc. Maybe it is just the price you are paying for that.

All those T_EX's 12

Hans Hagen
Hasselt NL
February 2020

This is about T_EX, the program that is used as part of the large suite of resources that make up what we call a 'T_EX distribution', which is used to typeset documents. There are many flavors of this program and all end with **tex**. But not everything in a distribution that ends with these three characters is a typesetting program. For instance, **latex** launches the a macro package L^AT_EX, code that feeds the program **tex** to do something useful. Other formats are Plain (no **tex** appended) or ConT_EXt (**tex** in the middle). Just take a look at the binary path of the T_EX distribution to get an idea. When you see **pdftex** it is the program, when you see **pdflatex** it is the macro package L^AT_EX using the pdfT_EX program. You won't find this for ConT_EXt as we don't use that model of mixing program names and macro package names.

Here I will discuss the programs, not the macro packages that use them. When you look at a complete T_EXLive installation, you will see many T_EX binaries. (I will use the verbatim names to indicate that we're talking of programs). Of course there is the original **tex**. Then there is its also official extended version **etex**, which is mostly known for adding some more primitives and more registers. There can be **aleph**, which is a stable variant of **omega** meant for handling more complex scripts. When pdf became popular the **pdftex** program popped up: this was the first T_EX engine that has a backend built in. Before that you always had to run an additional program to convert the native dvi output of T_EX into for instance PostScript. Much later, **xetex** showed up, that, like Omega, dealt with more complex scripts, but using recent font technologies. Eventually we saw **luatex** enter the landscape, an engine that opened up the internals with the Lua script subsystem; it was basically a follow up on **pdftex** and **aleph**.

The previous paragraph mentions a lot of variants and there are plenty more. For cjk and especially Japanese there are **ptex**, **eptex**, **uptex**, **euptex**. Parallel to **luatex** we have **luajittex** and **luahbtex**. As a follow up on the (presumed stable) **luatex** the ConT_EXt community now develops **luametatex**. A not yet mentioned side track is nts (New T_EX system), a rewrite of good old T_EX in Java, which in the end didn't take off and was never really used.

There are even more T_EX's and they came and went. There was **enctex** which added encoding support, there were **emtex** and **hugeemtex** that didn't add functionality but made more possible by removing some limits on memory and such; these were quite important. Then there were vendors of T_EX systems that came up with variants (some had extra capabilities), like **microtex**, **pctex**, **yandytex** and **vtex** but they never became part of the public effort.

For sure there are more, and I know this because not so long ago, when I cleaned up some of my archives, I found `eetex` (extended e- \TeX), and suddenly remembered that Taco Hoekwater and I indeed had experimented with some extensions that we had in mind but that never made it into e- \TeX . I had completely forgotten about it, probably because we moved on to $\text{Lua}\text{\TeX}$. It is the reason why I wrap this up here.

In parallel there have been some developments in the graphic counterparts. Knuts `metafont` program got a Lua enhanced cousin `mflua` while `metapost` (aka `mpost` or `mp`) became a library that is embedded in $\text{Lua}\text{\TeX}$ (and gets a follow up in $\text{LuaMeta}\text{\TeX}$). I will not discuss these here.

If we look back at all this, we need to keep in mind that originally \TeX was made by Don Knuth for typesetting his books. These are in English (although over time due to references he needed to handle different scripts than Latin, be it just snippets and not whole paragraphs). Much development of successors was the result of demands with respect to scripts other than Latin and languages other than English. Given the fact that (at least in my country) English seems to become more dominant (kids use it, universities switch to it) one can wonder if at some point the traditional engine can just serve us as well.

The original `tex` program was actually extended once: support for mixed usage of multiple languages became possible. But apart from that, the standard program has been pretty stable in terms of functionality. Of course, the parts that made the extension interface have seen changes but that was foreseeable. For instance, the file system hooks into the `kpse` library and one can execute programs via the `\write` command. Virtual font technology was also an extension but that didn't require a change in the program but involved postprocessing the `dvi` files.

The first major 'upgrade' was e- \TeX . For quite a while extensions were discussed but at some point the first version became available. For me, once $\text{pdf}\text{\TeX}$ incorporated these extensions, it became the default. So what did it bring? First of all we got more than 256 registers (counters, dimensions, etc.). Then there are some extra primitives, for instance `\protected` that permits the definition of unexpandable macros (although before that one could simulate it at the cost of some overhead) and convenient ways to test the existence of a macro with `\ifdefined` and `\ifcsname`. Although not strictly needed, one could use `\dimexpr` for expressions. A probably seldom used extension was the (paragraph bound) right to left typesetting. That actually is a less large extension than one might imagine: we just signal where the direction changes and the backend deals with the reverse flushing. It was mostly about convenience.

The Omega project (later followed up by Aleph) didn't provide the additional programming related primitives but made the use of wide fonts possible. It did extend the number of registers, just by bumping the limits. As a consequence it was much more demanding with respect to memory. The first time I heard of e- \TeX and Omega was at the 1995 euro \TeX meeting organized by the ntg and I was sort of surprised by the sometimes emotional clash between the supporters of these two variants. Actually it was the first time I became aware of \TeX politics in general, but that is another story. It was also

the time that I realized that practical discussions could be obscured by nitpicking about speaking the right terminology (token, node, primitive, expansion, gut, stomach, etc.) and that one could best keep silent about some issues.

The pdfTeX follow up had quite some impact: as mentioned it had a backend built in, but it also permitted hyperlinks and such by means of additional primitives. It added a couple more, for instance for generating random numbers. But it actually was a research project: the frontend was extended with so called character protrusion (which lets glyphs hang into the margin) and expansion (a way to make the output look better by scaling shapes horizontally). Both these extensions were integrated in the paragraph builder and are thereby extending core code. Adding some primitives to the macro processor is one thing, adapting a very fundamental property of the typesetting machinery is something else. Users could get excited: TeX renders a text even better (of course hardly anyone notices this, even TeX users, as experiments proved).

In the end Omega never took off, probably because there was never a really stable version and because at some time XeTeX showed up. This variant was first only available on Apple computers because it depends on third party libraries. Later, ports to other systems showed up. Using libraries is not specific for XeTeX. For instance pdfTeX uses them for embedding images. But, as that is actually a (backend) extension it is not critical. Using libraries in the frontend is more tricky as it adds a dependency and the whole idea about TeX was that is is independent. The fact that after a while XeTeX switched libraries is an indication of this dependency. But, if a user can live with that, it's okay. The same is true for (possibly changing) fonts provided by the operating system. Not all users care too strongly about long term compatibility. In fact, most users work on a document, and once finished store some pdf copy some place and then move on and forget about it.

It must be noted that where e-Tex has some limited right to left support, Omega supports more. That has some more impact on all kinds of calculations in the machinery because when one goes vertical the width is swapped with the height/depth and therefore the progression is calculated differently.

Naturally, in order to deal with scripts other than Latin, XeTeX did add some primitives. I must admit that I never looked into those, as ConTeXt only added support for wide fonts. Maybe these extensions were natural for L^AT_EX, but I never saw a reason to adapt the ConTeXt machinery to it, also because some pdfTeX features were lacking in XeTeX that ConTeXt assumed to be present (for the kind of usage it is meant for). But we can safely say that the impact of XeTeX was that the TeX community became aware that there were new font technologies that were taking over the existing ones used till now. One thing that is worth noticing is that XeTeX is still pretty much a traditional TeX engine: it does for instance OpenType math in a traditional TeX way. This is understandable as one realizes that the OpenType math standard was kind of fuzzy for quite a while. A consequence is that for instance the OpenType math fonts produced by the Gust foundation are a kind of hybrid. Later versions adopted some more pdfTeX features like expansion and protrusion.

I skip the Japanese \TeX engines because they serve a very specific audience and provide features for scripts that don't hyphenate but use specific spacing and line breaks by injecting glues and penalties. One should keep in mind that before Unicode all kinds of encodings were used for these scripts and the 256 limitations of traditional \TeX were not suited for that. Add to that demands for vertical typesetting and it will be clear that a specialized engine makes sense. It actually fits perfectly in the original idea that one could extend \TeX for any purpose. It is a typical example of where one can argue that users should switch to for instance $\text{Xe}\text{\TeX}$ or $\text{Lua}\text{\TeX}$ but these were not available and therefore there is no reason to ditch a good working system just because some new (yet unproven) alternative shows up a while later.

We now arrive at $\text{Lua}\text{\TeX}$. It started as an experiment in 2005 where a Lua interpreter was added to $\text{pdf}\text{\TeX}$. One could pipe data into the \TeX machinery and query some properties, like the values of registers. At some point the project sped up because Idris Hamid got involved. He was one of the few ConTeXt users who used Omega (which it actually did support to some extent) but he was not satisfied with the results. His oriental \TeX project helped pushing the $\text{Lua}\text{\TeX}$ project forward. The idea was that by opening up the internals of \TeX we could do things with fonts and paragraph building that were not possible before. The alternative, $\text{Xe}\text{\TeX}$ was not suitable for him as it was too bound to what the libraries provides (rendering then depends on what library gets used and what is possible at what time). But, dealing with scripts and fonts is just one aspect of $\text{Lua}\text{\TeX}$. For instance more primitives were added and the math machinery got an additional OpenType code path. Memory constraints were lifted and all became Unicode internally. Each stage in the typesetting process can be intercepted, overloaded, extended.

Where the $\text{e}\text{-}\text{\TeX}$ and Omega extensions were the result of many years of discussion, the $\text{pdf}\text{\TeX}$, $\text{Xe}\text{\TeX}$ and $\text{Lua}\text{\TeX}$ originate in practical demands. Very small development teams that made fast decisions made that possible.

Let's give some more examples of extensions in $\text{Lua}\text{\TeX}$. Because $\text{pdf}\text{\TeX}$ is the starting point there is protrusion and expansion, but these mechanisms have been promoted to core functionality. The same is true for embedding images and content reuse: these are now core features. This makes it possible to implement them more naturally and efficiently. All the backend related functionality (literal pdf, hyperlinks, etc) is now collected in a few extension primitives and the code is better isolated. This took a bit of effort but is in my opinion better. Support for directions comes from Omega and after consulting with its authors it was decided that only four made sense. Here we also promoted the directionality to core features instead of extensions. Because we wanted to serve Omega users too extended tfm fonts can be read, not that there are many of them, which fits nicely into the whole machinery going 32 instead of 8 bits. Instead of the $\text{e}\text{-}\text{\TeX}$ register model, where register numbers larger than 255 were implemented differently, we adopted the Omega model of just bumping 256 to 65536 (and of course, 16K would have been sufficient too but the additional memory it uses can be neglected compared to what other programs use and/or what resources users carry on their machines).

The modus operandi for extending \TeX is to take the original literate web sources and define change files. The $\text{pdf}\text{\TeX}$ program already deviated from that by using a monolithic source. But still Pascal is used for the body of core code. It gets translated to C before being compiled. In the $\text{Lua}\text{\TeX}$ project Taco Hoekwater took that converted code and laid the foundation for what became the original $\text{Lua}\text{\TeX}$ code base.

Some extensions relate to the fact that we have Lua and have access to \TeX 's internal node lists for manipulations. An example is the concept of attributes. By setting an attribute to a value, the current nodes (glyphs, kerns, glue, penalties, boxes, etc) get these as properties and one can query them at the Lua end. This basically permits variables to travel with nodes and act accordingly. One can for instance implement color support this way. Instead of injecting literal or special nodes that themselves can interfere we now can have information that does not interfere at all (apart from maybe some performance hit). I think that conceptually this is pretty nice.

At the Lua one has access to the \TeX internals but one can also use specific token scanners to fetch information from the input streams. In principle one can create new primitives this way. It is always a chicken-egg question what works better but the possibility is there. There are many such conceptual additions in $\text{Lua}\text{\TeX}$, which for sure makes it the most 'aggressive' extension of \TeX so far. One reason for these experiments and extensions is that Lua is such a nice and suitable language for this purpose.

Of course a fundamental part of $\text{Lua}\text{\TeX}$ is the embedded MetaPost library. For sure the fact that ConTeXt integrates MetaPost has been the main reason for that.

The ConTeXt macro package is well adapted to $\text{Lua}\text{\TeX}$ and the fact that its users are always willing to update made the development of $\text{Lua}\text{\TeX}$ possible. However, we are now in a stage that other macro packages use it so $\text{Lua}\text{\TeX}$ has entered a state where nothing more gets added. The $\text{L}\text{\TeX}$ macro package now also supports $\text{Lua}\text{\TeX}$, although it uses a variant that falls back on a library to deal with fonts (like $\text{X}\text{\TeX}$ does).

With $\text{Lua}\text{\TeX}$ being frozen (of course bugs will be fixed), further exploration and development is now moved to $\text{LuaMeta}\text{\TeX}$, again in the perspective of ConTeXt . I will not go into details apart from saying that is is a lightweight version of $\text{Lua}\text{\TeX}$. More is delegated to Lua, which already happened in ConTeXt anyway, but also some extra primitives were added, mostly to enable writing nicer looking code. However, a major aspect is that this program uses a lean and mean code base, is supposed to compile out of the box, and that sources will be an integral part of the ConTeXt code base, so that users are always in sync.

So, to summarize: we started with tex and moved on to etex and pdftex . At some point omega and xetex filled the Unicode and script gaps, but it now looks like luatex is becoming popular. Although luatex is the reference implementation, $\text{L}\text{\TeX}$ exclusively uses luahbtex , while ConTeXt has a version that targets at luametatex . In parallel, the $[\text{e}]$ $[\text{u}]$ $[\text{p}]$ tex engines fill the specific needs for Japanese users. In most cases, good old tex and less old etex are just shortcuts to pdftex which is compatible but has the pdf backend on board. That 8 bit engine is not only faster than the more recent

engines, but also suits quite well for a large audience, simply because for articles, thesis, etc. (written in a Latin script, most often English) it fits the bill well.

I deliberately didn't mention names and years as well as detailed pros and cons. A user should have the freedom to choose what suits best. I'm not sure how well \TeX would have evolved or will evolve in these days of polarized views on operating systems, changing popularity of languages, many (also open source) projects being set up to eventually be monetized. We live in a time where so called influencers play a role, where experience and age often matters less than being fancy or able to target audiences. Where something called a standard today is replaced quickly by a new one tomorrow. Where stability and long term usage of a program is only a valid argument for a few. Where one can read claims that one should use this or that because it is todays fashion instead of the older thing that was the actually the only way to achieve something at all a while ago. Where a presence on facebook, twitter, instagram, whatsapp, stack exchange is also an indication of being around at all. Where hits, likes, badges, bounties all play a role in competing and self promotion. Where today's standards are tomorrow's drawbacks. Where even in the \TeX community politics seem to creep in. Maybe you can best not tell what is your favorite \TeX engine because what is hip today makes you look out of place tomorrow.

Hidden treasures 13

Hans Hagen
Hasselt 2020
February 2020

At ConTeXt meetings we always find our moments to reflect on the interesting things that relate to TeX that we have run into. Among those we discussed were some of the historic treasures one can run into when one looks at source files. I will show examples from several domains in the ecosystem and we hereby invite the reader to come up with other interesting observations, not so much in order to criticize the fantastic open source efforts related to TeX, but just to indicate how decades of development and usage are reflected in the code base and usage, if only to make it part of the history of computing.

I start with the plain TeX format. At the top of that file we run into this:

```
% The following changes define internal codes as recommended
% in Appendix C of The TeXbook:
```

```
\mathcode`^\^^@="2201 % \cdot
\mathcode`^\^^A="3223 % \downarrow
\mathcode`^\^^B="010B % \alpha
\mathcode`^\^^C="010C % \beta
\mathcode`^\^^D="225E % \l and
\mathcode`^\^^E="023A % \l not
\mathcode`^\^^F="3232 % \in

...
\mathcode`^\^^Y="3221 % \rightarrow
\mathcode`^\^^Z="8000 % \neq
\mathcode`^\^^[="2205 % \diamond
\mathcode`^\^^\="3214 % \leq
\mathcode`^\^^\="3215 % \geq
\mathcode`^\^^\="3211 % \equiv
\mathcode`^\^^\="225F % \lor
```

This means that when you manage to key in one of these recommended character codes that in ascii sits below the space slot, you will get some math symbol, given that you are in math mode. Now, if you also consider that the plain TeX format is pretty compact and that no bytes are wasted,⁸ you might wonder what these lines do there. The answer is simple: there were keyboards out there that had these symbols. But, by the time TeX became popular, the dominance of the ibm keyboard let those memories fade away.

⁸ Such definitions don't take additional space in the format file.

This is just Don's personal touch I guess. Of course the question remains if the sources of TAOCP contain these characters.

There is another interesting hack in the plain \TeX file, one that actually, when I first looked at the file, didn't immediately made sense to me.

```
\font\preloaded=cmti9
\font\preloaded=cmti8
\font\preloaded=cmti7

\let\preloaded=\undefined
```

What happens here is that a bunch of fonts get defined and they all use the same name. Then eventually that name gets nilled. The reason that these definitions are there is that when \TeX dumps a format file, the information that comes from those fonts is embedded to (dimensions, ligatures, kerns, parameters and math related) data. It is an indication that in those days it was more efficient to have them preloaded (that is why they use that name) than loading them at runtime. The fonts are loaded but you can only access them when you define them again! Of course nowadays that makes less sense, especially because storage is fast and operating systems do a nice job at caching files in memory so that successive runs have font files available already.

Talking of fonts, one of the things a new \TeX user will notice and also one of the things users love to brag about is ligatures. If you run the `tftopl` program on a file like `cmr10.tfm` you will get a verbose representation of the font. Here are some lines:

```
(LABEL C f)  (LIG C i 0 14) (LIG C f 0 13) (LIG C l 0 15)
(LABEL O 13) (LIG C i 0 16) (LIG C l 0 17)
(LABEL C `)  (LIG C ` C \`)
(LABEL C ')  (LIG C ' C ")
(LABEL C -)  (LIG C - C {)
(LABEL C {)  (LIG C - C |)
(LABEL C !)  (LIG C ` C <)
(LABEL C ?)  (LIG C ` C >)
```

The `C` is followed by an ascii representation and the `)` by the position in the font `O` (a number) or `C` (a character). So, consider the first two lines to be a puzzle: they define the `fi`, `ff`, `fl` ligatures as well as the `ffi` and `ffl` ones. Do you see how ligatures are chained?

But anyway, what do these other lines do there? It looks like ```` becomes the character in the backslash slot and `''` the one in the double quote. Keep in mind that \TeX treats the backslash special and when you want it, it will be taken from elsewhere. But still, these two ligatures look familiar: they point to slots that have the left and right double quotes.⁹ They are not really ligatures but abuse the ligature mechanism to achieve a similar effect. The last four lines are the most interesting: these are ligatures that

⁹ \TeX never assumed this and encourages users to use the quotation macros. Those ```quotes''` look horrible in a source anyway.

(probably) no \TeX user ever uses or encounters. They are again something from the past. Also, changes are low that you mistakenly enter these sequences and the follow up Latin Modern fonts don't have them anyway.

Actually, if you look at the MetaFont and MetaPost sources you can find lines like these (here we took from `mp.w` in the Lua \TeX repository):

```

@ @<Put each...@=
mp_primitive (mp, "=:", mp_lig_kern_token, 0);
@:=:_}{\.{=:} primitive@;
mp_primitive (mp, "=:|", mp_lig_kern_token, 1);
@:=:_}{\.{=:}\char'174 primitive@;
mp_primitive (mp, "=:|>", mp_lig_kern_token, 5);
@:=:_}{\.{=:}\char'174> primitive@;
mp_primitive (mp, "|=:", mp_lig_kern_token, 2);
@:=:_}{\.{=\char'174=:} primitive@;
mp_primitive (mp, "|=:>", mp_lig_kern_token, 6);
@:=:_}{\.{=\char'174=:>} primitive@;
mp_primitive (mp, "|=:|", mp_lig_kern_token, 3);
@:=:_}{\.{=\char'174=:}\char'174 primitive@;
mp_primitive (mp, "|=:|>", mp_lig_kern_token, 7);
@:=:_}{\.{=\char'174=:}\char'174> primitive@;
mp_primitive (mp, "|=:|>>", mp_lig_kern_token, 11);
@:=:_}{\.{=\char'174=:}\char'174>>} primitive@;

```

I won't explain what happens there (as I would have to reread the relevant sections of \TeX The Program) but the magic is in the special sequences: `=: =:| =:|>|=:| =:|>|=:|>>`. Similar sequences are used in some font related files. I bet that most MetaPost users never entered these as they relate to defining ligatures for fonts. Most users know that combining a `f` and `i` gives a `fi` but there are other ways to combine too. One can praise today's capabilities of OpenType ligature building but \TeX was not stupid either! But these options were never really used and this treasure will stay hidden. Actually, to come back to a previous remark about abusing the ligature mechanism: OpenType fonts are just as sloppy as \TeX with the quotes: there a ligature is just a name for a multiple-to-one mapping which is not always the same as a ligature.

But there are even more surprises with fonts. When Alan Braslau and I redid the bibliography subsystem of Con \TeX t with help from Lua, I wrote a converter in that language. I actually did that the way I normally do: look at a file (in this case a bib \TeX file) and write a parser from scratch. However, at some point we wondered how exactly strings got concatenated so I decided to locate the source and look at it there. When I scrolled down I noticed a peculiar section:

```

@^character set dependencies@>
@^system dependencies@>
Now we initialize the system-dependent |char_width| array, for which
|space| is the only |white_space| character given a nonzero printing

```

width. The widths here are taken from Stanford's June-'87 \$cmr10\$-font and represent hundredths of a point (rounded), but since they're used only for relative comparisons, the units have no meaning.

```
@d ss_width = 500          {character '@'31|'s width in the $cmr10$ font}
@d ae_width = 722          {character '@'32|'s width in the $cmr10$ font}
@d oe_width = 778          {character '@'33|'s width in the $cmr10$ font}
@d upper_ae_width = 903   {character '@'35|'s width in the $cmr10$ font}
@d upper_oe_width = 1014  {character '@'36|'s width in the $cmr10$ font}

@<Set initial values of key variables@>=
for i:=0 to @'177 do char_width[i] := 0;
@#
char_width[@'40] := 278;
char_width[@'41] := 278;
char_width[@'42] := 500;
char_width[@'43] := 833;
char_width[@'44] := 500;
char_width[@'45] := 833;
```

Do you see what happens here? There are hard coded font metrics in there! As far as I can tell, these are used in order to guess the width of the margin for references. Of course that won't work well in practice, simply because fonts differ. But given that the majority of documents that need references are using Computer Modern fonts, it actually might work well, especially with Plain \TeX because that is also hardwired for 10pt fonts. Personally I'd go for a multipass analysis (or maybe would have had $\text{bib}\text{\TeX}$ produce a list of those labels for the purpose of analysis but for sure at that time any extra pass was costly in terms of performance). That code stays around of course. It makes for some nice deduction by historians in the future.

I bet that one can also find weird or unexpected code in ConTeXt, and definitely on the machines of \TeX users all around the world. For instance, now that most people use utf8 all those encoding related hacks have become history. On the other hand, as history tends to cycle, bitmap symbolic fonts suddenly can look modern in a time when emoji are often bitmaps. We should guard our treasures.

Don't use \TeX ! 14

Occasionally I run into a web post that involves \TeX and it is sometimes surprising what nonsense one can read. Now, I normally don't care that much but there are cases where I feel that a comment is justified. Partly because I'm one of the developers, but also because I'm involved in user groups.

In this particular case the title of a (small) blog post was "*Why I do not like \TeX !*" and the website announced itself ambitiously as '*DIGITAL TYPOGRAPHY NEWS*'. Normally I assume that in such a case it is a general site about typesetting and that the author has not much experience or insight in the already ancient \TeX typesetting system. However, the url is:

eutypon.gr/e-blog/index.php/2021/02/13/why-i-do-not-like-luatex/

which happens to be the Greek User Groups portal. So why do I feel the need to reflect on this? Why do I even care? The answer is simple: because user groups should inform their (potential) users correctly. Another reason is that I'm involved in the program that is disliked, and yet another one is that there is a suggestion that language support is bad in \TeX , while actually hyphenation patterns are very well maintained by Mojca and Arthur who are also actively involved in the community around the mentioned engine.

Let's start with the title. For sure one might not like a specific program, but when it involves one of the mainstream \TeX engines, it should at least be clear that it's a personal opinion. Because no name is mentioned, we can assume that this is the opinion of the Greek user group as a whole. The text starts with "*Most people speak with good words about \TeX .*" and the '*most*' in that sentence sort of puts the author in a small group, which should trigger using a bit more careful title. Now I know a couple of users who use \TeX (with ConTeXt) for typesetting Greek, and we can assume that they are among the people who speak those good words: typesetting Greek just works.

More good news is that "*They seem to like things it can do that no other \TeX derived system can do.*" This might invite potential users to take a closer look at the system, especially because we already know that most people are positive. In 2021 one should keep in mind that, although the \TeX engine is around for more than a decade, the level of support can differ per macro package which is why pdfTeX is still the most widespread used \TeX variant: much \TeX usage relates to writing (scientific articles) in English so one doesn't really need an Unicode engine. I always say: don't change a good working workflow if you have no reason; use what makes you feel comfortable. Only use \TeX if you have a reason. There is plenty of good and positive advice to be given.

With "*Personally, I do not care about these features but yesterday a friend told me that he wanted to write something in Greek with \TeX .*" the author steps over his or her personal rejection of the engine and enters the help-a-friend mode. "*And what's the catch, one may*

Why I do not like luatex!

Most people speak with good words about luatex. They seem to like things it can do that no other TeX derived system can do. Personally, I do not care about these features but yesterday a friend told me that he wanted to write something in Greek with luatex. And what's the catch, one may ask. The problem is that luatex does not load any hyphenation patterns but the default ones. So one needs to load them. In TeX one uses a command like the following one

```
\language\l@monogreek%
```

where `\l@monogreek` is numerical value assigned to each language contained in the format. This is well documented in the TeXbook. Now despite the fact that I spent a few hours searching for information on how to load specific hyphenation patterns, I could not find anything! Moreover, I could not find any information on how one loads a lua package (i.e., some external lua package that is available in the TeX installation). People know that they can load a LaTeX package with the `\usepackage` command but I have no information on how to load lua code. Practically, this means that if one is not part of the inner circle of luatex developers, then she cannot really know what is really going on. And this is exactly the reason why I do not like luatex.

Published February 13, 2021
By [admin](#)

Categorized as [Uncategorized](#)

eutypon.gr/e-blog/index.php/2021/02/13/why-i-do-not-like-luatex/

ask. The problem is that *luLaTeX* does not load any hyphenation patterns but the default ones. So one needs to load them." I'm not sure why this is a catch. It actually is a feature. One drawback of the traditional \TeX engines is that one needs to preload the hyphenation patterns. Before memory was bumped, that often meant creating format files for a subset of languages, and when memory became plenty it meant preloading dozens of patterns by default. The good news is that in all these cases the macro package takes care of that. In the case of *LuaTeX* no patterns need to be preloaded so it might even be that *LATeX* doesn't have any preloaded but, not being a user, I didn't check that.

This all makes the next sentence puzzling: "*In TeX one uses a command like the following one : \language\l@monogreek, where \l@monogreek is numerical value assigned to each language contained in the format.*" Now, I'm no expert on *LATeX* but I'm pretty sure that the \@ sign is not a letter by default. I'm also pretty sure that there is some high level interface to enable a language, and in the case of *LuaTeX* being used that mechanism will load the patterns runtime. I bet it will also deal with making sure other language specific properties are set. Therefore the "*This is well documented in the TeXbook.*" is somewhat weird: original \TeX only had one language and later versions could deal with more, but plain \TeX has no \l@monogreek command. It doesn't sound like the best advice to me.

Just to be sure, I unpacked all the archives in the most recent $\text{\TeX}Live$ dvd and grepped for that command in tex and sty files and surprise: in the *LATeX* specific style file $/tex/xelatex/xgreek/xgreek.sty$ there is a line $\text{\language\l@monogreek\else\HyphenRules{monogreek}\fi}$ which to me looks way to low level for common users to figure out, let alone that it's a file for \XeTeX so bound to a specific engine. Further grepping for $\{greek\}$ gave hits for *LATeX*'s *babel* an there are Greek files under the *polyglossia* directory so I bet that Arthur (who once told me he was responsible for languages) deals with Greek there. Even I, as a *ConTeXt* user who never use *LATeX* and only know some things by rumor (like the fact that there is something like polyglossia at all) could help a new user with some suggestions of where to look, just by googling for a solution. But explicitly using the \language primitive is not one of them. Okay, in *ConTeXt* the \language [greek] command does something useful, but we're not talking about that package here, if only because it relates to *LuaTeX* development, which as we will see later is a kind of inner circle.

So, picking up on the blog post, in an attempt to get Greek working in *LATeX* the author got online but "*Now despite the fact that I spent a few hours searching for information on how to load specific hyphenation patterns, I could not find anything!*" It might have helped to search for *lualatex greek* because that gives plenty of hits. And maybe there are even manuals out there that explain which of the packages in the \TeX tree to load in order to get it working. Maybe searching *ctan* or $\text{\TeX}Live$ helps too. Maybe other user groups have experts who can help out. No matter what you run into, I don't think that the average user expects to find a recipe for installing and invoking patterns. Just for the record, the *LuaTeX* manual has a whole chapter on language support, but again, users can safely assume that the macro package that they use hides those details. Actually, if users were supposed to load patterns using a unique id, they are likely to end up in the modern Greek versus ancient Greek, as well as Greek mixed with English or other

languages situations. That demands some more in depth knowledge to deal with, in any macro package and with any engine. You can add a bit of Unicode and utf-8 or encodings in the mix too. Suggesting to consult the $\text{\TeX}book$ is even a bit dangerous because one then also ends up in an eight bit universe where font encodings play a role, while LuaTeX is an Unicode engine that expects utf and uses OpenType fonts. And, while languages seem to be a problem for the author and his/her friend, fonts seem to be an easy deal. In my experience it's more likely that a user runs into font issues because modern fonts operate on multiple axis: script, language and features.

Maybe the confusion (or at that time accumulated frustration) is best summarized by “*Moreover, I could not find any information on how one loads a lua package (i.e., some external lua package that is available in the TeX installation).*” Well, again I’m sure that one can find some information on $\text{L}\text{\TeX}$ support sites but as I already said: language support is so basic in a macro package that users can use some simple command to enable their favorite one. So, when “*People know that they can load a LaTeX package with the \usepackage command but I have no information on how to load lua code.*” the first part is what matters: Lua files are often part of a package and thereby they get loaded by the package, also because often stand-alone usage makes not much sense.

It is absolutely no problem if someone doesn’t like (or maybe even hates) LuaTeX , but it’s a different matter when we end up in disinformation, and even worse in comments that smell like conspiracy: there is an inner circle of LuaTeX developers and “*Practically, this means that if one is not part of the inner circle of luaTeX developers, then she cannot really know what is really going on.*” Really? Is this how user groups educate their users? There are manuals written, plenty of articles published, active mailing lists, presentations given, and there is support on platforms like Stack Exchange. And most of that (the development of LuaTeX included) is done by volunteers in their spare time, for free. Of course the groups of core developers are small but that is true for any development. History (in the \TeX community) has demonstrated that this is the only way to make progress at all, simply because there are too many different views on matters, and also because the time of volunteers is limited. It is the end result what counts and when that is properly embedded in the community all is fine. So we have some different engines like \TeX , pdfTeX , LuaTeX , etc., different macro packages, specialized engines like those dealing with large cjk fonts, all serving a different audience from the same ecosystem. Are these all secretive inner circles with bad intentions to confuse users?

The blog post ends with “*And this is exactly the reason why I do not like luatex .*” to which I can only comment that I already long ago decided not to waste any time on users who in their comments sound like they were forced to use a \TeX system (and seem to dislike it, so probably are better off with Microsoft Word, but nevertheless like to bark against some specific \TeX tree), who complain about manuals not realizing that their own contributions might be rather minimalistic, maybe even counter productive, or possibly of not much use to potential users anyway. I also ignore those who love to brag about the many bugs, any small number suits that criterium, without ever mentioning

how bugged their own stuff is, etc. If your ego grows by disregarding something you don't even use, it's fine for me.

So why do I bother writing this? Because I think it is a very bad move and signal of a user group to mix personal dislike, whatever the reason is, with informing and educating users. If a group is that frustrated with developments, it should resolve itself. On the other hand, it fits well in how todays communication works: everyone is a specialist, which get confirmed by the fact that many publish (also on topics they should stay away from) on the web without fact checking, and where likes and page hits are interpreted as a confirmation of one's expertise. Even for the \TeX community there seems to be no escaping from this.

The objectives of \TeX user groups shift, simply because users can find information and help online instead of at meetings and in journals. The physical \TeX distributions get replaced by fast downloads but they are definitely under control of able packagers. Maybe a new task of user groups is to act as guardian against disinformation. Of course one then has to run into these nonsense blogs (or comments on forums) and such but that can partly be solved by a mechanism where readers can report this. A user group can then try to make its own information better. However, we have a problem when user groups themselves are the source of disinformation. I see no easy way out of this. We can only hope that such a port drowns in the ocean of information that is already out there to confuse users. In the end a good and able \TeX friend is all you need to get going, right? The blog post leaves it open if the Greek text ever got typeset well. If not, there's always ConTeXt to consider, but then one eventually ends up with LuaMetaTeX which might work on the author as another "*rode lap op een stier*" as we say in Dutch.

15.1 Introduction

Recently a couple of cordless phones that I use gave up as soon as I used them for a minute or so. The first time that happened I figured that after all these years the batteries had gone bad and after some testing I decided to replace them. I got some of these high end batteries that discharge slowly and store a lot of power. Within a year they were dead too. Then I went for the more regular and cheaper ones, again with a lot of capacity. And yes, these also gave up, that is: only in the phones that were hardly used. The batteries lasted longer in phones that were discharged by usage daily.

When I went out for new batteries I was asked if I needed them for cordless phones and, surprise, was given special ones that actually stored less but were guaranteed to work for at least 6 years. The package explicitly mentioned use in cordless phones. So here performance doesn't come with the most high end ones, based on specifications that impress.

This is also true for computers that are used to process \TeX documents. More cores amount to much accumulated processing power but for a single core \TeX process, a few fast cores are more relevant than plenty slower ones that run in parallel. More memory helps but compared to other processes \TeX actually doesn't use that much memory. And disk speed matters but less so when the operating system caches files. What does play a role are cpu caches because \TeX is very memory intense and processing is not concentrated in a few functions. But a large cache shared among many (busy) cores makes for a less impressive performance.

So what matters really? In the next sections we will explore a few points of view. It's not some advertisement for a specific engine, but much more about putting it into perspective (as one can run into ridiculous arguments on the web). It is not only the hardware and software that matters but also how one uses it.

15.2 The engine

There are various ways to compare engines and each has its own characteristics. The $\text{pdf}\text{\TeX}$ engine is closest to the original. It directly produces the output which can give it an edge. It is eight bit and therefore uses small fonts and internally all that is related to fonts and characters is also small. This means that there is little overhead in typesetting a paragraph: hyphenation, ligature building and kerning are interwoven and perform well.

The $\text{Xe}\text{\TeX}$ engine supports wide fonts and Unicode and therefore can be seen as 32 bit. I never looked into the code so I can't tell how far that goes but performance is

definitely less than pdfTeX. The rendering of text is delegated to a library (there were some changes in that along its development) which is less efficient than the built in pdfTeX route. But it is also more powerful.

The LuaTeX engine is mostly 32 bit and delegates non standard font handling to Lua which comes with a performance penalty but also adds a lot of flexibility. Also, the fact that one can call out to Lua in many places makes that one can not really blame the engine for performance hits. The fact that hyphenation, ligature building and kerning is split comes at a small price too. We have larger nodes so compared to pdfTeX more memory is used and accessed. Some mechanisms are actually more efficient, like font expansion and protrusion.

The LuaMetaTeX engine lacks a font loader (but it does have the traditional renderer on board) and it has no backend. So even more is delegated to Lua, which in turn makes this the slowest of the lot. And, again more data travels with nodes. In some modes of operation much more calculations take place. However, because it has an enriched macro processor, additional primitives, and plenty deep down '*improvements*' it can perform better than LuaTeX (and even LuajitTeX, the LuaTeX version with a faster but limited Lua virtual machine). And as with LuaTeX, there are usage patterns that make it faster than pdfTeX.

So, in general the order of performance is pdfTeX, XeTeX, LuajitTeX (kind of obsolete), LuaTeX, LuaMetaTeX. But then, how come that ConTeXt users never complain about performance? The reasons is simple: performance is quite okay and as it is relative to what one does, a user will accept a drop in performance when more has to be done. When we moved on from LuaTeX to LuaMetaTeX there definitely was a drop in performance, simply because of the Lua backend. Because upgrading happened in small (but continuous) steps, right from the start the new engine was good enough to be used in production which is why most users switched to LMTX as soon as became clear that this is where the progress is made.

There were no real complaints about the upto 15% initial performance drop which indicates that for most users it doesn't matter that much. As the engine evolved we could gain some back and now LuaMetaTeX ends up between pdfTeX and LuaTeX and in many modern scenarios even comes out first. The fact that in the meantime we can be much faster than LuaTeX did get noticed (when asked). However, as development takes years updating a machine in the meantime puts discussions about performance in a different (causality) perspective anyway.

15.3 The coding

Performance can increase when native engine features are used instead of complex macros that have to work around limitations. It can also decrease when new features are used that add complex functionality. And when an engine extends existing functionality that is likely to come at a price. So where LuaMetaTeX provides a more rich

programming environment, it also had a more complex par builder, page builder, insert, mark and adjust handling, plenty of extra character, rule and box features and all of that definitely adds some overhead. Quite often a gain in simplicity (nicer and more efficient macros) compensate the more complex features. That is because on the average the engine doesn't do that much (tens of thousands of the same) complex macro expansion and also doesn't demand that much complex low level typesetting. A gain here is often compensated by a loss there. This is one reason why during the years LuaMetaTeX could sustain a decent performance. Personally I don't accept a drop in performance easily which is why in practice most mechanism, even when extended, probably perform better but I'm not going to prove that observation.

One important reason why ConTeXt LMTX with LuaMetaTeX is faster than its ancestors is that we got rid of some intermediate programming layers. Most users have never seen the auxiliary macros or implementation details but plenty were used in MkII and MkIV. Of course we kept them because often they are nicer than many lines of primitive code, but only a few (and less in the future) are used in the core. Examples are multi step macros (that pick up arguments) that became single step and complex if tests that became inline native tests. Because ConTeXt always had a high level of abstraction consistency of the interface also makes that we don't need many helpers. When some features (like for instance box manipulation) got extended one could expect a performance hit due to more extensive optional keyword scanning in the engine but that was compensated by improved scanners. The same is true for scanning numbers and dimensions. So, more functionality doesn't always come at a price.

To summarize this: although the engine went a bit more '*cisc*' than *risc* the macro package went more '*risc*'. It reminds me a bit of the end of the previous century when there was much talk of fourth generation languages, something on top of the normal languages. In the end it were scripting languages that became the fashion while traditional languages like C remained relatively stable and unchanged for implementing them (and more). A similar observation can be made for ConTeXt itself. Whenever some new feature gets added to an existing mechanism I try to not cripple performance and thanks to the way ConTeXt is set up it works out okay.

Let's look at an example. In MkII we can compare two '*strings*' with the macro *doifelse*. Its definition is as follows:

```
\long\def\doifelse#1#2%
{\let\donottest\donotprocess
 \edef\!!stringa{#1}%
 \edef\!!stringb{#2}%
 \let\donottest\doprocess
 \ifx\!!stringa\!!stringb
   \expandafter\firstoftwoarguments
 \else
   \expandafter\secondoftwoarguments
 \fi}
```

This macro takes two arguments that gets expanded inside two helpers that we then compare with a primitive `\ifx`. Depending on the outcome we expand one of the two following arguments but first we get rid of the interfering `\else` and `\fi`. The pushing and popping of `\donottest` takes care of protection of unwanted expansion in an `\edef`. Many functional macros are what we call protected: then expand in two steps depending on the embedded `\donottest` macro. Think of (simplified):

```
\def\realfoo{something is done here}
\def\usedfoo{\donottest\realfoo}
```

Normally `\donottest` is doing nothing so `\realfoo` gets expanded but there are cases where we (for instance) `\let` it be `\string` which then serializes the macro. This is something that happens when writing to the multi pass data file. It can also be used for overloading, for instance in the backend or when converting something. This protection against expansion has always been a ConTeXt feature, which in turn made it pretty robust in multi pass scenarios, but it definitely came with performance penalty.

When pdfTeX got the e-TeX extensions we could use the `\protected` prefix to replace this trickery. That means that MkII will use a different definition of `\doifelse` when that primitive is known:

```
\long\def\doifelse#1#2%
{\edef\!{stringa}{#1}%
 \edef\!{stringb}{#2}%
 \ifx\!{stringa}\!{stringb}
   \expandafter\firstoftwoarguments
 \else
   \expandafter\secondoftwoarguments
 \fi}
```

This works okay because we now do this:

```
\protected\def\usedfoo{something is done here}
```

The `\doifelse` helper itself is not protected in MkII (non e-TeX mode) It would be a performance hit. I won't bore the reader with the tricks needed to do the opposite, that is: expand a protected macro. It is seldom needed anyway.

The MkIV definition used with LuaTeX is not much different, only the `\long` prefix is missing. That one is needed when one wants `#1` and/or `#2` to be tolerant with respect to embedded `\par` equivalents. In LuaMetaTeX we can disable that check and in ConTeXt all macros are thereby `\long`. Users won't notice because in ConTeXt most macros were always defined the long way; we also suppress `\outer` errors.

```
\protected\def\doifelse#1#2%
{\edef\m_syst_string_one{#1}%
 \edef\m_syst_string_two{#2}%
 \ifx\m_syst_string_one\m_syst_string_two
```

```

  \expandafter\firstoftwoarguments
\else
  \expandafter\secondoftwoarguments
\fi}

```

Implementation wise a macro, once scanned and stored, carries the long property in its command code so that has overhead. However because LuaTeX is compatible we cannot make all normal macros long by default when $\text{\textbackslash suppresslongerror}$ is used. Therefore checking for an argument running into a $\text{\textbackslash par}$ is still checked but the message is suppressed based on the setting of the mentioned parameter. Performance wise, not using $\text{\textbackslash long}$ comes a the cost of checking a parameter which means an additional memory access and comparison. Unless we otherwise gain something in the engine it comes at a cost. In LuaMetaTeX the $\text{\textbackslash long}$ and $\text{\textbackslash outer}$ prefixes are ignored. Even better, protected macros are also implemented a bit more efficiently.

In the end the definition of $\text{\textbackslash doifelse}$ in LMTX looks a bit different:

```

\permanent\protected\def\doifelse#1#2%
{\iftok{#1}{#2}%
  \expandafter\firstoftwoarguments
\else
  \expandafter\secondoftwoarguments
\fi}

```

The $\text{\textbackslash permanent}$ prefix flags this macro as such. Depending on the value of $\text{\textbackslash overloadmode}$ a redefinition is permitted, comes with a warning or results in a fatal error. Of course this comes at a price when we define macros or values of quantities but this is rather well compensated by all kind of improvements in handling macros: defining, expansion, saving and restoring, etc.

More interesting is the use of $\text{\textbackslash iftok}$ here. It saves us defining two helper macros. Of course the content still needs to be expanded before comparison but we no longer have various macro management overhead. In scenarios where we don't need to jump over the $\text{\textbackslash else}$ or \textbackslash fi we can use this test in place which saves passing two arguments and grabbing one argument later on. Actually, grabbing is also different, compare:

```

\def\firstoftwoarguments #1#2{#1} % MkII and MkIV
\permanent\def\firstoftwoarguments #1#-{#1} % MkXL aka LMTX

```

```

\def\secondoftwoarguments#1#2{#1} % MkII and MkIV
\permanent\def\secondoftwoarguments#-#1{#1} % MkXL aka LMTX

```

In the case of LuaMetaTeX the $\text{\textbackslash #--}$ makes that we don't even bother to store the argument as it is ignored. Where \textbackslash #0 does the same it also increments the argument counter which is why here even the second arguments has number 1. Now, if this more efficient? Sure, but how often does it really happen? The engine still needs to scan (which comes at a cost) but we save on temporary token list storage. Because \TeX is so fast already,

measuring only shows differences when one has many (and here a real lot) iterations. However, all these small bits add up which is what we've seen in 2022 in ConTeXt: it is the reason why we are now faster than MkIV with LuaTeX, even with more functionality in the engine.

I can probably write hundreds of pages in explaining what was added, changed, made more flexible and what side effects it had/has on performance but I bet no one is really interested in that. In fact, the previous exploration is just a side effect of a question that triggered it, so maybe future questions will trigger more explanations. It anyhow demonstrates what I meant when I said that LuaMetaTeX is meant to be leaner and meaner. Of course the code base and binary is smaller but that also gets compensated by more functionality. It also means that we can make the ConTeXt code base nicer because for me a good looking source (which of course is subjective) is pretty important.

15.4 Compatibility

There are non ConTeXt users who seem to love to stress that successive versions of ConTeXt are incompatible. Other claims are that it is developed in a commercial setting. While it is true that there are changes and it is also true that ConTeXt is used in commercial settings, it is not that different from other open source projects. The majority of the code is written without compensation and it is offered without advertisements or request for support. It is true that when we can render better, it will be done. But the user interfaces only change when there is a reason and there are few cases where some functionality became obsolete, think of input and font encodings. Most such changes directly relate to the engine: in pdfTeX and MkII we emulate utf-8 while in LuaTeX it comes natively. In pdfTeX eight bit (Type1) fonts are used while LuaTeX adds support for OpenType. Other macro packages support that by additional packages while ConTeXt has it integrated. That is why the system evolves over time.

Just as users adapt to (yearly) operating system interfaces, mobile phones, all kinds of hardware, cars, clothing, media and so on, the ConTeXt users have no problem adapting to an evolving TeX ecosystem. I guess claims about changes (being a disadvantage) can only point to a lack of development elsewhere. The main reason for mentioning this is that when ConTeXt users move on to newer engines, the older ones are seldom used. So, few users compare a LMTX run with one using pdfTeX or LuaTeX. They naturally expect LuaMetaTeX to perform well and maybe even to perform better over time. They just don't complain. And unless one hacks (overloads) system macros compatibility is not really an issue. What can be an issue is that updates and adaptations to a newer engine come with bugs but those are solved.

So, the fact that we compare incompatible engines with likely different low level macro implementations of otherwise stable features of a macro package makes comparison hard. For instance, maybe there are speedups possible in frozen MkII, although it is unlikely, which makes that it might even perform better than reported. In a similar fashion, the fact that OpenType is more demanding for sure makes that LuaTeX rendering is slower than pdfTeX. It anyhow makes a discussion about performance within

and between macro packages even more ridiculous. Just don't buy those claims and/or ask on the ConTeXt mailing list for clarification.

15.5 The job

So, say that we now have an efficient and powerful engine and a matching macro package. Does that make all jobs faster? For sure, the ones that I use as benchmark run much smoother. The 360 page LuaMetaTeX manual runs in less than 8.4 seconds on a Dell Precision laptop with (mobile) Intel(R) Xeon(R) CPU E3-1505M v6 @ 3.00GHz, 2TB fast Samsung pro SSD, and 48 GB of memory, running Windows 10. The MetaFun manual with many more pages and thousands of MetaPost graphics needs a bit more than 12 seconds. So you don't hear me complain. This chapter takes 7.5 seconds plus 0.5 is for the runner, not enough time to get coffee.

Nowadays I tend to measure performance in terms of pages per second, because in the end that is what users experience. For me more important are the gains for my colleague who processes documents of 400 pages from hundreds of small xml files with multiple graphics per page. Given different output variants a lot of processing takes place, so there a gain from 20 pages per second to 25 pages per second is welcome. Anyway, here are a few measurements of a *simple* test suite per January 7, 2023. We use this as test text:

```
\def\Knuth{%%
Thus, I came to the conclusion that the designer of a new system
must not only be the implementer and first large||scale user; the
designer should also write the first user manual.
\par
The separation of any of these four components would have hurt
\TeX\ significantly. If I had not participated fully in all these
activities, literally hundreds of improvements would never have
been made, because I would never have thought of them or perceived
why they were important.
\par
But a system cannot be successful if it is too strongly influenced
by a single person. Once the initial design is complete and fairly
robust, the real test begins as people with many different
viewpoints undertake their own experiments.
}
```

Now keep in mind that these are simple examples. On more complex documents the LuaMetaTeX engine with LMTX is relatively faster: think xml, plenty MetaPost, complex tables, advanced math, dozens of fonts in combination with the new compact font mode.

The tests themselves are simple: we switch fonts (because fonts bring overhead), we add some color (because we use different methods), we process some graphics (to show

what embedding MetaPost brings), we do some tables (because that can be stressful). Each sample is run 50, 500 or 1000 times, and each set is run a couple of times so that we compensate for caching and fluctuating system load. The tests are more about signaling a trend than about absolute numbers. For what it's worth, I used a Lua script to run the samples.

When you run an experiment that measures performance, keep in mind that performance not only depends on the engine, but also on for instance logging. When I run the ConTeXt test suite it takes 1250 seconds if the console takes the full screen on a 2560 by 1600 display and 30 seconds more on a 3840 by 2160 display and it even depends on how large the font is set. On the 1920 by 1200 monitor I get to 1230. Of course these times change when we add more to the test suite so it's always a momentary measurement.

Similar differences can be observed when running in an editor. A good test is making a ConTeXt format: 2.2 seconds goes down to below 1.8 when the output is piped to a file. On a decent 2023 desktop those times are probably half but I don't have one at hand.

sample 1, number of runs: 2

```
\starttext
  \dorecurse {%-s} {
    \Knuth
    \par
  }
\stoptext
```

engine	50	500	1000
pdftex	0.63	0.83	1.07
luatex	0.95	1.86	2.94
luametatex	0.61	1.49	2.48

sample 2, number of runs: 2

```
\starttext
  \dorecurse {%-s} {
    \tf \Knuth \bf \Knuth
    \it \Knuth \bs \Knuth
    \par
  }
\stoptext
```

engine	50	500	1000
pdftex	0.70	1.73	2.80
luatex	1.37	5.37	9.92
luametatex	1.04	5.06	9.73

sample 3, number of runs: 2

```
\starttext
  \dorecurse {%-s} {
    \tf \Knuth \it knuth \bf \Knuth \bs knuth
    \it \Knuth \tf knuth \bs \Knuth \bf knuth
    \par
  }
\stoptext
```

engine	50	500	1000
pdftex	0.71	1.81	2.98
luatex	1.41	5.84	10.77
luametatex	1.05	5.71	10.60

sample 4, number of runs: 2

```
\setupcolors[state=start]
\starttext
  \dorecurse {%-s} {
    {\red \tf \Knuth \green \it knuth}
    {\red \bf \Knuth \green \bs knuth}
    {\red \it \Knuth \green \tf knuth}
    {\red \bs \Knuth \green \bf knuth}
    \par
  }
\stoptext
```

engine	50	500	1000
pdftex	0.73	1.91	3.64
luatex	1.39	5.82	12.58
luametatex	1.07	5.57	11.85

sample 5, number of runs: 2

```
\starttext
  \dorecurse {%-s} {
    \null \page
  }
\stoptext
```

engine	50	500	1000
pdftex	0.62	1.12	1.68
luatex	0.90	1.39	1.98
luametatex	0.58	0.99	1.46

sample 6, number of runs: 2

```
\starttext
  \dorecurse {%-s} {
    %% nothing
  }
\stoptext
```

engine	50	500	1000
pdftex	0.55	0.54	0.56
luatex	0.79	0.81	0.82
luametatex	0.54	0.52	0.53

sample 7, number of runs: 2

```
\starttext
  \dontleavehmode
  \dorecurse {%-s} {
    \framed[width=1cm,height=1cm,offset=2mm]{x}
  }
\stoptext
```

engine	50	500	1000
pdftex	0.58	0.65	0.71
luatex	0.84	0.96	1.08
luametatex	0.54	0.62	0.72

sample 8, number of runs: 2

```
\starttext
  \dontleavehmode
  \dorecurse {%-s} {
    \framed
      [width=1cm,height=1cm,offset=2mm,
       foregroundstyle=bold,foregroundcolor=red,
       background=color,backgroundcolor=green]
    {x}
  }
\stoptext
```

engine	50	500	1000
pdftex	0.59	0.70	0.83
luatex	0.87	1.00	1.17
luametatex	0.55	0.66	0.78

sample 9, number of runs: 2

```
\starttext
  \ifdef\permanent\else\def\BC{\NC\bf}\fi
  \dontleavehmode
  \dorecurse {%-s} {
    \starttabulate[|||||]
      \NC test \BC test \NC test \NC test \NC \NR
      \NC test \BC test \NC test \NC test \NC \NR
      \NC test \BC test \NC test \NC test \NC \NR
      \NC test \BC test \NC test \NC test \NC \NR
    \stoptabulate
  }
\stoptext
```

engine	50	500	1000
pdftex	0.62	1.15	1.71
luatex	0.94	1.84	2.86
luametatex	0.60	1.19	1.88

sample 10, number of runs: 2

```
\starttext
  \dontleavehmode
  \dorecurse {%-s} {
    \startMPcode
      fill fullcircle scaled 1cm withcolor red ;
      fill fullsquare scaled 1cm withcolor green ;
    \stopMPcode
    \space
  }
\stoptext
```

engine	50	500	1000
pdftex	5.73	50.98	102.10
luatex	0.93	1.07	1.30
luametatex	0.57	0.71	0.86

15.6 Final words

Whenever I run into (or get send) remarks of (especially non ConTeXt) users suggesting that LuaTeX is much slower than pdfTeX or that LuaMetaTeX seems much faster than LuaTeX, one really has to keep in mind that this is not always true. Among the questions to be asked are “*What engine do you use?*”, “*Which macro package do you use?*”, “*How well is your style set up?*”, “*How complex is the document?*”, “*Is your own additional code efficient?*”,

“Do you use engine and macro package features the right way?” and of course *“What do you compare with?”*, *“What do you expect and why?”*, *“Do you actually know what goes on deep down?”*. An embarrassing one can be *“Do you have an idea what is involved in fulfilling your request given that we use a flexible adaptive macro language?”*. Much probably these questions not get answered properly.

Another thing to make clear is that when someone claims for instance that ConTeXt LMTX is fast because of LuaMetaTeX, or that LuaMetaTeX is much faster than LuaTeX, a healthy suspicion should kick in: does that someone really knows what happens and matters? The previous numbers do show differences for simple cases but we’re often not talking of differences that can be used as an excuse for insufficient coding. In the end it is all about the experience: does performance feel in tune with expectations. Which is not to say that I will make ConTeXt and LuaMetaTeX faster because after all there are usage scenarios where one has to process tens of thousands of documents with a reasonable amount of time, on regular infrastructure, and of course with as little as possible energy consumption.

If pdfTeX suits your purpose, there is no need to move to LuaTeX. As with rechargeable batteries in cordless phones a higher capacity can make things worse. If LuaTeX fits the bill, don’t dream about using LuaMetaTeX instead because it will half runtime because the adaptations needed in the macro package (like adding a backend) might actually slow it down. Moores law doesn’t apply to TeX engines and macro packages and you might get disappointed. Accept that the choice you made for a macro package can come with a price.

Quite often it is rather easy to debunk complaints and claims which makes one wonder why claims about perceived or potential are made at all. But then, I’m accustomed to weird remarks and conclusions about ConTeXt as a macro package, or for that matter LuaTeX (as it originates in the ConTeXt community) even by people who should know better. Hopefully the above invites to being more careful.

16.1 Introduction

When working on a \TeX macro package for decades one can hardly avoid dealing with math; after all \TeX is pretty much about math. When this wonderful typesetting infrastructure was written it was all about quality and how to make your documents look nice. And for sure, Don Knuth's documents look nice, also because he pays a lot of attention to the “*fine points of math typesetting*”.

The constraints of those times (like hardware, compilers, fonts, and for sure also time) made \TeX into what it is: eight bit character sets, eight bit fonts, eight bit hyphenation patterns, efficient memory usage and therefore carrying around as little as possible. It all makes sense. But one needs to pay attention.¹⁰

Math typesetting is actually a sort of separated process in the engine: unprocessed lists go in and after some juggling a list of assembled boxes, glyphs, glues and penalties come out. I will not go into detail about that and only mention that in LuaMetaTeX we extended all this to be a bit more flexible and controllable, something that has been driven by the fact that we need to support Unicode fonts. This is all part of a related effort to move from eight bit ‘*everything*’ to Unicode ‘*everywhere*’.

Now, one can say a lot about Unicode but the main advantage is that it tries to cover ‘*all*’ characters ever encountered, including scripts (used in languages) that are long gone, as well as these little pictures that people like to see on the web: emojis. One can safely say that Unicode simplifies mixing languages and scripts, and thereby makes \TeX macro packages less complex. On the other hand, Unicode (or more precisely, related wide) fonts makes all kind of features possible and thereby add a complication.

So, how about math? When Don Knuth gave us \TeX he also gave us fonts and there are plenty symbols in these fonts. But, as mathematicians seem to love variations on symbols soon more fonts arrived, most noticeably those from the ams that also added some more alphabets: mathematicians also love to render the shapes of letters differently. In order to access these glyphs names were invented that also sometimes suggested that there was some order in the matter. And, for some reason these names got aliases and soon we had a huge list of often obscure and inconsistent macro names. It didn't take long for a little mess and confusion to creep in.

It has been said that the verbose \TeX math ascii input format is also a way for mathematicians to communicate, just because many use the same tool to render the formulas.

¹⁰ And that is what Mikael Sundqvist and I have been doing a lot since we started upgrading math in ConTeXt in combination with enhancing the math engine in LuaMetaTeX . The story here is a byproduct of our explorations and very much a combined effort.

Of course that gets obscured when one starts to add additional macros. It gets even more tricky once we start talking ‘*standard*’ as in “*L^AT_EX is the standard*”. That has for instance resulted in browsers interpreting \TeX like input without using \TeX (so how about expansion?). It has also sort of put \TeX into the range of possible word processing systems, which in turn leads to these MS Word versus Google docs versus L^AT_EX debates that can get rather nasty and unrealistic when it comes to discussing usage and quality. Interestingly, MS Word now has reasonable math, to some extent modelled after \TeX . It has some verbose \TeX like (but constrained) input and would do well for probably mostly people who occasionally have to inject some math. There were also attempts by the people at Microsoft to normalize the input but we leave that aside now.

However, because we now do have all these symbols and because source code editors make them accessible and show them there is a good chance that users will inject them, if only by cut and paste, so we do have to deal with that. This automatically puts us in the position that we need to deal with different meanings for the same symbol, which in turn might demand different spacing, penalties and such. In the end it is users that drive all this, not publishers; they don’t really care and out-source typesetting anyway. We’re not aware of any research and development being done and I suppose we would have noticed because after all we’re involved in developing Lua \TeX . It is one of the engines that does OpenType and Unicode math and no publisher or supplier ever took serious interest in it. From our perspective what users do is visible, everything else is hidden behind corporate curtains. And this is why nowadays we only need to care about users (mainly authors).

Back to typesetting. For a long time all went well: one could typeset documents that looked good. Okay, not all looked good because not everyone paid attention to details, and the more the web evolved the more patching cut’n’paste of bad examples made its way into documents, but let’s not start talking quality here. But then came Unicode and a while later people started talking about accessibility, cutting and pasting and more. In the meantime there had been developments like MathML and OpenMath that tried to structure and organize formulas in a more symbolic way.¹¹

In the meantime the \TeX community had lost the edge on fonts, and OpenType math was invented by Microsoft and implemented in MS Word before a substantial number of \TeX users understood what was happening. They had it coming. To a large extend one can say the same about math in Unicode. Where a Greek capital ‘A’ is seen as different from a Latin capital ‘A’, even when they often have the same shape, a math italic variable ‘ h ’ was made synonym to ‘*Planck constant*’, as if the letters used in math had no meaning at all. We’ll see that a wide hat is an extensible character of zero width combining hat accent, which makes for curious handling of the initial character. There is more granularity in some symbols, especially popular symbols like slashes and bars, than in

¹¹ It probably went unnoticed that Con \TeX t always supported rendering MathML, and as such had to deal with all the weird aspects (read: way it was used). Although one is not supposed to directly edit MathML we work with authors who are quite happy to do that simply because they code the documents in xml because there is a need for high quality pdf as well as html output and a Con \TeX t based workflow can handle the xml well. We’re talking of large volumes here (mostly for basically free school math).

letters. It is as if the math community didn't care much about how the letters (variables) were communicated and perceived but were picky about the slope of slashes. It seems more of a visual world, which might actually be the reason structured input never really took off. Maybe \TeX ies just love the mix of characters, commands, spacing directives. Maybe they just love to reposition and space these glyphs to suit all kind of curious non-standard math rendering.

All this makes it pretty hard to communicate meaning, and it is just one of the examples where the \TeX community, for as far involved, failed to make a strong case. Our personal opinion is that no one really cared because in the \TeX community it is all about rendering. The fact that we use math to communicate only gained attention when accessibility became hot and by then it was too late. Efforts like OpenMath started ambitious and in the end basically failed. Coding in xml using MathML isn't much better and one always had to adapt to the latest fashion. Also, once plenty code shows up bugs become features. Browser support came and went and came back. Simplified input using for instance AsciiMath started indeed simple but quickly became a (somewhat inconsistent) mess. What we see here is the same as everything web (and computer languages): we can do better, we start some project, then move on, and we end up with half-way abandoned results. The development cycles are short, results have to be achieved fast, there is no time (or interest) for iterating and refactoring. The word '*standard*' and mantra '*everyone should use this*' are quite popular.

So where does that leave us with \TeX ? Well, with a mess. Decades of various efforts have not brought us a coherent system of organizing symbols and properties, made us end up with inconsistencies, made users revert to hacks, didn't make math easily transferable and complicates rendering. Personally we find it sort of strange that we spend time on for instance tagging and accessibility before we get these math alphabets and shared math specific symbols sorted out. If we cannot make good arguments for that (math being a script on its own with semantics and such) we waste energy and are pulling a dead horse. What puzzles us most is that one would expect mathematicians to be able to come up with strong arguments for a structured approach. But maybe it was simply the fact that \TeX math typesetting was pretty much driven by large commercial publishers and those providing services for them: the first category doesn't invest in these matters and even less today, and the second category makes money from sorting out the mess, so why get rid of it. Who knows. For us, it means that any complain about these matters deserves the same answer: the \TeX community created this mess, so it has to live with it. And the bad thing is: bugs and work-arounds eventually become features and then one is supposed to conform, even if deep down one knows better. It doesn't help that the community is proud of what it can render and has built itself a reputation that all is good.

So why this criticism? Why not just abandon \TeX ? The answer is simple: \TeX is quite okay and cannot be blamed for where we are now. We need to think of solutions and in that respect the Con \TeX t users are lucky! They have always been told not to use this macro package for math because there are other standards and because publishers want L \TeX (even if they just let the manuscripts be recoded). That means that we

don't really need to care much about the past. Those who use ConTeXt can benefit from the compatibility we have anyway but also move forward to more structured and consistent math. It is in this perspective that we will discuss some more details next so that eventually we can draw some conclusions. The end goal is to have an additional layer of grouping math symbols that permits consistent high quality rendering in a mixed input environment.

16.2 Molecules

Before we go into details about some characters, we spend some word on the rendering. The building blocks of a formula are atoms and internally the term nucleus is used for what we have without scripts. The simple sequence $1 + x$ will result in a linked list of three atoms with three nuclei. In x^2 the x is the nucleus. Atoms can have scripts: prescripts, postscripts and a prime. The majority of Unicode math characters become such atoms (nuclei and scripts) and they get a class property that determines their spacing, but that is not part of the Unicode specification. From the upcoming sections it will be clear that when we classify we don't get that much help from MathML or even the TeX community either.

In addition to these atoms the LuaMetaTeX engine (which builds upon TeX) has what we can call molecules. There are several types: fractions, accents, fences, radicals. This distinction is to some extent present in Unicode: plenty of fraction related slashes, all kind of accents, vertical delimiters that can be made from snippets and act as fences, and a radical symbol. In MathML we see similar constructs but there in practice quite often operators need to be interpreted in a way that can distinguish between atoms and molecules. That is partly a side effect of applications that generate MathML. And as usual with standards pushed upon the world without years of exploration the confusion became part of the norm and will stay.

In the TeX engine over and under delimiters are implemented on top of radicals (using the same noad, the wrapper node for yet unprocessed math) but they have different code paths. Basically we have vertically fenced material and just like fractions have left and right fences as part of the concept (for binomials) the radical has a sort of left fence too. You can also wonder why we need accent noads while we support other delimiters with radicals. This organization mostly relates to subtypes and classes (and likely some limitations of the past) that have related spacing properties, but we can think if a generic structure noad and meaningful subtypes. However, that is not what we get so let's be more precise:

Fractions: these stack two atoms (or molecules) and separate them by a visible or phantom rule, or in LuaMetaTeX by a delimiter. They can have a left and right fence which originates in them also suitable for binomials. You may wonder why we don't use regular fences here. One reason we can think of is that when you fence something, you have an open and close class at the edges while with a fenced fraction the whole still is fraction. In LuaMetaTeX we can tweak classes at the edges but in regular TeX there are fewer classes, so there constructs become ordinary or inner.

Accents: these put something on top of or below an atom (or molecule) and are driven by characters. The accent related commands take an integer (traditional) or three integers (extended) and it is this expected input that drives it. However, they are treated like delimiters. In traditional \TeX a delimiter is defined by two characters: the direct unscaled one, and when not found a second one drives the lookup from wider variants and eventually an extensible character. Accents just have the second one, which probably relates to the fact that the text ones that would be the starting point make no sense. It is this ‘*looking*’ for a single code point that makes that accents are not merged with the more general radical command space. Another reason is that accents deal a bit different with spacing and italic correction so even if we could merge, it would be more confusing in the end.

Fences: these come in pairs with optional middle ones. The reason for pairing is that they need to get the same size. That means that before we construct them the atom or molecule that they fence has to be analyzed. It also makes the result a construct of its own, although in LuaMetaTeX we can unpack that result so that it can be broken across lines. In practice that was never an issue because in a running text unscaled fences are used (just atoms with open and close classes assigned) but as soon as one goes to multi-line displays formulas things become more hairy. The related commands expect delimiters (the two part character definitions) but in the meantime are also happy with a single one because in the end OpenType math has all in one font.

Radicals: originally this only concerned roots but because they are basically wrappers we also use them for content that gets a delimiter above, below or both. In that sense the term radical can also be interpreted as ‘*extreme*’, more than a carrot looking symbol. The related commands take one or more delimiters (or character) because we support left as well as right delimiters connected by a rule, so in the end radicals evolved into a construct with delimiters of all kind. So, the unique property of radicals is that the fences assume a cooperation between one or more glyphs and a rule. In ConTeXt we support actuarian hooks as radicals that are used for annuity expressions, otherwise the Unicode symbols is useless and the MathML construct complex.

So, where accents take numbers as delimiter specification, fences, fractions and radicals take specific math quantities or just letters. This makes that we will not merge these into one scanner and handler even if they all use the same (large) node to store and carry around their properties. Also, it has some charm to keep the original \TeX distinctions. After all, it’s not like Unicode, MathML or OpenType math fonts have brought some new insights: in the end they all draw from \TeX and they way it’s done there.

16.3 Symbols

There are plenty of symbols in Unicode. When we try to get an idea how we ended up with that set we’re surprised that not much seems to be known about it. There are references to iso standards, usage by specific organizations (like those dealing with patents), there are references to lists of publishers. In personal communications with people involved it becomes clear that the criterion that some symbols really has to be

used somewhere doesn't apply to these math symbols. There are bizarre specimens that we cannot locate anywhere. They are often assigned the '*relation*' property which for \TeX is a safe bet because binary and relations get similar spacing, but binary makes an exception when it sits at the front. The fact that relation spacing is used can even obscure the fact that some characters have zero width properties; the results just look somewhat bad and one can always blame the font or renderer and adding some thin spacing is accepted behavior. So one can make the argument that because \TeX was the main renderer of math, a safe bet was better than a confusing and unproven-by-usage assignment to some category.

In \TeX some symbols have multiple names, even when they have the same class. This indicates the wish for meaning at one end but shape at the other, and once a name has been assigned it sticks. It would be interesting to know how mathematicians see formulas: if one puts `\bars` around a variable does one see "*bar x bar*" or "*the modulus of x*", and how is translation to audio to be performed?

One important aspect of using any symbol in \TeX , or basically any typesetting system that deals with math, is that the spacing depends on the meaning. Now, in the perspective of Unicode meaning is somewhat diffuse. A Latin capital 'A' related to 'a' is not the same as a Greek capital 'A' that relates to 'α'. So, from the shape one cannot beforehand deduce what is meant, but when copying it the Unicode will expose the meaning. This is not the case in math: although many symbols have one meaning only, there are also plenty that can mean different things and the (\TeX) math community has not been able to make a strong case for providing different slots. Maybe the reason was that there already was a tradition of using commands that then relate a shape to a class that then results in appropriate spacing. Maybe it is also assumed that an article or book starts by explaining what a specific symbol means in that particular context. But that doesn't help much for copying. It also doesn't help with direct Unicode input. The way out for this last problem is that in ConTeXt we will add additional properties to characters that then can communicate the class and thereby control the spacing. Although we initially did that at the Lua end we now use the lightweight dictionary feature of the engine: a property, group, slot model. The main reason is that we foresee that at some point we might have to add property based rendering to the engine, and this opens up that possibility. Ever since we started with LuaTeX and MkIV we have used the character database (in Lua format) to store most properties so that we have all in one place.

For figuring out the properties we can look at how traditionally symbols got multiple commands associated, how MathML looks at it, what Unicode reveals and what we find in fonts. It is a bit of jungle out there so for sure we have to make decisions ourselves. We next turn to that exploration.

16.4 Slashes

The definition on the Wikipedia page [1] of slashes is as follows:

"The slash is an oblique slanting line punctuation mark /. Once used to mark periods and commas, the slash is now used to represent exclusive or inclusive or, division and

fractions, and as a date separator. It is called a solidus in Unicode, is also known as an oblique stroke, and has several other historical or technical names including oblique and virgule.

The page then has a very detailed description on how slashes are used in text, mathematics, computing, currency, dates, numbering, linguistic transcriptions, line breaks, abbreviations, proofreading, fiction, libraries, addresses, poetry, music, sports, and text messages. It is a pretty good and detailed page which also gives a nice summary of usage in math.

In mathematics, we use the slash (a forward leaning bar) for fractions, division, and quotient of set. Examples of fractions are $\frac{1}{2}$ but also $\%$ sits in this category.

U+0002F	/	this is the official solidus
U+02044	/	the mathematical fraction slash
U+02215	/	the mathematical division slash
U+02571	/	a diagonal box drawing line
U+029F8	/	the mathematical big solidus
U+0FF0F	?	a full width solidus
U+1F67C	?	the very heavy solidus

The Stix fonts have the first five, the rest is not there, so we can safely assume that they are not used in math. That brings us to the question that, say that the other ones are used, how does the user access them? In the editor they often look pretty much the same. For \TeX ies the answer is easy: you use a command. But as we already mentioned, there we enter a real fuzzy area: these commands either describe a shape or they communicate a meaning, at least, in an ideal world. Sometimes wrapping in a macro helps, like $\$\\vfrac{1}{2}\$$.

In the document that explains Unicode math there is a section “*Fraction Slash and Other Diagonals*”. Even if we limit ourselves to the forward leaning slashes it looks like we need to include exotic symbols, as the empty set symbol with an left arrow on top: **U+29B4** a circle with left pointing arrow on top, that doesn’t show up in most math fonts but Stix has it \emptyset . We quote:

“**U+2044 / FRACTION SLASH** is typically used to build up simple skewed fractions in running text. It applies to immediately adjacent sequences of decimal digits, that is, to spans of characters with the General Category property value **Nd**. For example, $\frac{1}{2}$ should be displayed as $\frac{1}{2}$. In ordinary plain text, any character other than a digit delimits the numerator or denominator. So $5 \frac{1}{2}$ should be displayed as $5\frac{1}{2}$ since a space follows the 5 . In general mathematical use, a more versatile method for layout of fractions is needed (see, for example, Section 2.1 of [UnicodeMath]), however parsers of mathematical texts should be prepared to handle **FRACTION SLASH** when it is received from other sources. **U+27CB MATHEMATICAL RISING DIAGONAL** and **U+27CD MATHEMATICAL FALLING DIAGONAL** are mathematical symbols for specific uses, to be distinguished from the more widely used solidi and reverse solidi operators as well as from nonmathematical diagonals.”

In TeX there is no parsing going on: we just get sequences of atoms and the inter atom spacing applies. Curly braced arguments are used to communicate units that needs to be treated a while. As side note: where for some scripts there are special characters that tell where something (state) starts and ends this is not available for math, which makes it impossible to mark a sequence of characters as being something math. The whole repertoire of pre-composed fractions and super- and subscripted Unicode symbols are not to be used in math.

Most documents that somehow relate to or (partially) originate in TeX can be rather fuzzy, so we can read here:

“U+27CB corresponds to the L^AT_EX entity `\diagup` and U+27CD to `\diagdown`. Their glyphs are invariably drawn with 45° and 135° slopes, respectively, instead of the more upright slants typical for the solidi operators. The diagonals are also to be distinguished from the two box drawing characters U+2571 and U+2572. While in some fonts those characters may be drawn with 45° and 135° slopes, respectively, they are not intended to be used as mathematical symbols. One usage recorded for U+27CB and U+27CD is in the notation for spaces of double cosets.”

So, it is the angles that math users should translate into meaning which I guess is natural for them. From the above we cannot deduce if we should take them into account in a macro package.

The MathML specification [3] keeps it abstract and talks about division without mentioning the rendering. In content MathML we have:

```
divide = element divide { CommonAtt, DefEncAtt, empty}
```

and the suggested rendering (from an example) is a slash.

In the chapter “*Characters, Entities and Fonts*” there is mentioning of:

“There is one more case where combining characters turn up naturally in mathematical markup. Some relations have associated negations, such as U+226F [NOT GREATER-THAN] for the negation of U+003E [GREATER-THAN SIGN]. The glyph for U+226F [NOT GREATER-THAN] is usually just that for U+003E [GREATER-THAN SIGN] with a slash through it. Thus it could also be expressed by U+003E-U+0338 making use of the combining slash U+0338 [COMBINING LONG SOLIDUS OVERLAY]. That is true of 25 other characters in common enough mathematical use to merit their own Unicode code points. In the other direction there are 31 character entity names listed in [Entities] which are to be expressed using U+0338 [COMBINING LONG SOLIDUS OVERLAY].”

A curious note is this:

“For special purposes, one may need a symbol which does not have a Unicode representation. In these cases one may use the `mglyph` element for direct access to a glyph as an image, or (in some systems) from a font that uses a non-Unicode encoding. All MathML

*token elements accept characters in their content and also accept an **mglyph** there. Beware, however, that use of **mglyph** to access a font is deprecated and the mechanism may not work in all systems. The **mglyph** element should always supply a useful alternative representation in its alt attribute.”*

At some point we experimented with very precise positioned html from \TeX (read: ConTeXt) and that worked very well: the rendering was exactly the same as pdf but then suddenly it was no longer possible to access glyphs from fonts. The assumption had become that one should feed text into the font rendering machinery and use OpenType features to access specific shapes, which of course is a fragile approach (the libraries and logic keep evolving, and the most robust access is simply by index, or by glyph name if present, assuming that one uses the font that was meant to be used). So, how the MathML **glyph** element is supposed to work out well is not clear. Anyway, as we want nicely typeset math we don’t care that much if features present in LuaMeta \TeX and ConTeXt are unique and cannot be reproduced otherwise.

In **mathclass.txt** [4] which is “*not formally part of the Unicode Character Database at this time*” we see a classification:

U+0002F	binary
U+02044	binary
U+02215	binary
U+02571	not mentioned
U+029F8	n-ary or large operator, often takes limits
U+0FF0F	not mentioned
U+1F67C	not mentioned

So, in the end we can focus on the four that are mentioned, and we will do that with the above in mind as well as what is common in the \TeX world. We will look at usage, classification (groups) and classes.

Unfortunately this sort of mess also results in a mess in fonts. For instance when we checked out the difference between U+002F and U+2044 we found that in the fonts produced by the \TeX Gyre project both have proper dimensions (and look the same), so they can be used stand alone, but also as delimiters. In Cambria the dimensions are okay but only U+2044 has extensible characters. In ConTeXt we have defined \slash to use that slot but when you test Lucida and Stix2 the results are disappointing: In Lucida the width of U+2044 makes it unusable (it looks bad anyway), and in Stix2 it is a bit wider so in the end it even becomes fuzzy what to recommend as fix: quarter width, half width or full width. Defining \slash as any of them gives at some point an issue so in the end we just patch the font in the goodie file: we make them the same and make sure they have extensible characters. After all, chances are slim that this will ever be fixed. In that respect a newer engine doesn’t change the problem: we need to handle it in the macro package, but at least that can be done a bit more natural.¹²

¹² In principle, we can support the goodies in the generic font handler, but we think it makes no sense because it also relates to the way math is handled in general and supporting a wide range of different applications can only cripple the code, let along that agreeing on matters can be hard.

16.5 Bars

Again we start with the Wikipedia page, this time the one dedicated to bars [5]. The page starts with mathematics so that suggests that the (initial) author is familiar with usage in that field: if we cut and paste the itemized list we even get TeX math (sort of). Examples of usage are: absolute value, cardinality, conditional probability, determinant, distance, divisibility, function evaluation, length, norm, order, restriction, set-builder notation, the Sheffer stroke in logic, subtraction, but also “*A vertical bar can be used to separate variables from fixed parameters in a function, or in the notation for elliptic integrals*”.

Among the objectives of our exploration are grouping symbols in sets that represent related meanings and usage. Within these groups we can fine tune with classes but that is more geared at rendering. Although currently users enter specific usage of symbols with the same shape (or even Unicode) with commands we can imagine them entering the ‘*real*’ characters and in that case we need some automatic class assignment based on a group (or set of groups). The Wikipedia page mentions that in physics “*The vertical bar is used in bra–ket notation in quantum physics*”. It then goes on about usage in computing, phonetics and literature. This ordering is different from the slashes, but okay.

The page then makes a distinction between solid and broken bars and there is some interesting history behind that, which relates to typewriters, terminals and printers in the perspective of distinction and indeed we noticed that on our keyboard the broken bar is still used, even if the rendering is solid. The page ends with the Unicode bars and entities. We mention most:

U+007C		a single vertical line
U+00A6		a single broken line
U+2016		a double vertical line (norms)
U+2223		divides
U+2225		parallel lines
U+2502		a vertical box drawing line
U+FF5C	?	a fullwidth vertical line

Given the mentioned wide range of usage it will be clear bars that can be confusing and are pretty overloaded. We’re not aware of broken bars being used in math, so we ignore these.

The Unicode math draft talks of ‘*vertical lines*’ and distinguishes two series, delimiters:

U+007C		single vertical lines
U+2016		double vertical lines
U+2980		triple vertical lines

and operators:

U+2223		divides (single line)
U+2225		parallel (double lines)

U+2AF4 ||| binary relation (tripple lines)

U+2AFC ||| s large triple operator

Watch the triples: these are not (yet) in the Wikipedia summary. Rightfully there is a remark that the official Unicode descriptions use **BAR** and **LINE** but **TeXies** can't complain about that, can they? After all, they also use these terms mixed.

The delimiters sit at the edges but sometimes also in the middle. The operators are between other elements and the document states that they also should grow. And is it mentioned that spacing depends on usage. The large triple is an n-ary operator but as usual with math symbols the user (reader) has to guess what that actually means.

It is actually unfortunate that the fences have no left, middle and right variant. Even if these render the same it would make life easier and consistency with other fences is also worth something. One wonders how it would have looked if accessibility demands had kicked in earlier. The Unicode **mathclass.txt** [4] provides:

U+007C fence (unpaired delimiter)

U+2016 fence (unpaired delimiter)

U+2980 fence (unpaired delimiter)

We assume that the unpaired qualification is actually an indication that usage as what in **TeX** is called '*middle*' is okay. The operators are classified as:

U+2223 relation

U+2225 relation

U+2AF4 binary

U+2AFC large n-ary

The main problem with bars in **TeX** is that there is no distinction between a left and right bar which makes it impossible to use them directly as fences. One can consider this to be an omission to Unicode math because shape rules over meaning. So anyway, this is something that a macro package has to deal with. If needed these can get a class on their own in which case we can define atom spacing rules that deal with them ending up left or right. In Unicode there are signals that deal with bidirectional text, so we see no reason why there shouldn't be similar provisions for math.

16.6 Hyphens and Dashes

This section applies to text and math as both are riddled with horizontal lines: easy to scratch in wood, chisel in stone or draw on paper symbols. We limit ourselves to the straight ones, but similar observations can be made for curved ones.

Wikipedia distinguishes hyphens, minus, and dashes so there are multiple pages dedicated to this. The page about minus mentions that there are three usages (somewhat rephrased):

- It is used as subtraction operator and therefore a binary operator that indicates the operation of subtraction.
- It can be function whose value for any real or complex argument is the additive inverse of that argument.
- It can serve as a prefix of a numeric constant. When it is placed immediately before an unsigned numeral, the combination names a negative number, the additive inverse of the positive number that the numeral would otherwise name.

The functional variant is how content MathML sees it: you apply a minus operator to something, singular or multiple. We were surprised to see that there is a distinctive rendering suggested, something we have argued for at several occasions (mostly \TeX meetings):

*"In many contexts, it does not matter whether the second or the third of these usages is intended: $\text{‐}5$ is the same number. When it is important to distinguish them, a raised minus sign ‐ is sometimes used for negative constants, as in elementary education, the programming language *apl*, and some early graphing calculators."*

Unfortunately that distinction was not recognized by the \TeX community at large which (we guess) is why we don't see it in Unicode, which on the other hand has plenty dashes as we will see soon.

The page mentions usage in indicating blood types and music, which is a nice detail. It also mentions usage in computing, including regular expressions and in physics and chemistry indicating charge. It lists these codes for minus symbols:

U+002D hyphen minus
 U+2212 minus
 U+FE63 small hyphen minus
 U+FF0D full width hyphen minus

The page also mentions the commercial minus ‐ (see also [7]) and division sign ÷ (see also [8]) and we think these should be supported in math mode simply because they can be part of (even simple text style) formulas.

The fact that we use the hyphen as minus and expect it to render as a wider dash like shape is something that related to math mode in \TeX speak. In text mode we expect it to be seen as hyphenation related indicator. We won't go into details about automated hyphenation and explicit hyphens in text mode but here are the hyphens as mentioned on the hyphen specific Wikipedia page:

U+002D hyphen minus
 U+00AD soft hyphen
 U+2010 hyphen
 U+2011 non breaking hyphen

You might wonder why we mention text variants here and one reason is that we actually might need to provide a catch for the last two: maybe when a user copies these from a

document (when rendered at all) we need to treat them as the simple hyphen minus and just remap them to the math minus when in math mode. Below, we will discuss dashes, and although these are also meant for text, a reason for exploring these can be found in the fact that \TeX users like to decorate the content in unexpected ways and lines (or rules) fit into that. The Wikipedia pages go into some details about the hyphens being used in compounds and there can be some confusion about whether to use endashes or hyphens for that. We're pretty sure that typesetting wars have been fought over that. Usage as pre- and suffixes definitely is worth noting (and we use them as such in this sentence).

We leave out all the other usages and see what there is to tell about related symbols. The Wikipedia page about dashes is an extensive one. It starts out with the distinction between figure dash (U+2012: $-$), endash (U+2013: $-$), emdash (U+2014: $—$) and horizontal bar (U+2015: $—$). Of these a \TeX ie will for sure recognize the endash and emdash. The hyphen is not a dash but if you look at \TeX input that double or triple hyphens get ligatured into en- and emdashes! The only certainty one has is that the endash is often half the width of an emdash. Also, the width of the emdash is often the same as the font size.

One reason why a language subsystem of a \TeX macro package is complex is that it has to deal with cultural aspects and the usage as well as spacing around all these dashes can differ. When trying to support that a macro writer soon finds out that one user of language X can tell you the rules are done this way, and a while later you get a mail from another user who claims that in language X the rules are done that way. Word processing and dominance of English probably adds to the confusion. The same is true for quotes, but math doesn't need these, so we skip them. Now wait, you will say: does math use these dashes? Users probably will mix them in but more important is that the width of these dashes also has associated skips: `\enspace` and `\emspace` or `\quad` and these one definitely see users mix into math.

The figure dash has the same width as digits which makes them useful in tables. In the fonts that come with \TeX it is the reverse: the digits have the same width and that width matches the endash. There is no habit of using the figuredash, but we might need to change that. After all, we now have the fonts! We do need to deal with the figure dash because users might mix math and text in tables, and although you can find plenty of badly typeset by \TeX tables, this is no excuse for using a mix of minus and figure dash in inconsistent ways.

The Wikipedia page mentions the usage of the endash: as connector, as compound hyphen, and as sentence interrupter. Now the one that needs some attention is the second one. In Dutch, we can combine words in many ways and for educational purposes adding a compound dash makes sense. However, because the weight of the hyphen and endash in \TeX fonts is rather incompatible, in ConTeXt we use(d) fakes: two overlapping hyphens. Another complication is that one has to wrap that in a discretionary node in order to make the hyphenator happy, but that is now delegated to the engine that can be configured to see certain characters as valid hyphenation points. Although

we support discretionaries in math this doesn't relate to dashes but to pluses and minuses and such. The engine supports explicit discretionaries but can also automatically repeat symbols that are set up as repeatable across lines. We're not sure if users actually use en- and emdashes in math mode, but one can occasionally run into examples (on the web) where special effects are achieved in curious ways.¹³

It is worth pointing out that Wikipedia discusses “*Ranges of values*” and this is something we need to investigate in the perspective of math! Strictly spoken that is a text thing, but . . . Among the many observed and suggested patterns we note that among TeXies using the endash as itemize symbols is also popular.

Usage of the emdash is related to the use of parenthesis or colons, so it is more a kind of punctuation. It can also be used as an interrupt and again it is a candidate for an itemize symbol. There is of course a \TeX thing there: lack of text symbols made for a rather mixed usage of math and text symbols in itemizations. For instance a dotted one uses the well visible math dot instead of the often hardly visible text dot that simply was not present in \TeX fonts, so our eyes got accustomed to the bolder ones. It is one of the reasons why a \TeX macro package load a math font even when no math is used. Over the years in \TeX math and text symbols have been mixed in various ways, also a side effect if the limited amount of characters in text fonts and the abundance of them in math mode, even if most are only accessible by name. We need to deal with that historic mix.

The page rightfully mentions that `\TeX` has no horizontal bar, also known as ‘*quotation dash*’, used for dialogues in some languages. We should make a note then that it might be good to see if we have to reconfigure the sub-sentence presets to match that expectation. The proposed hack [MPS: where?](#) for a missing symbol is somewhat curious:

x \hbox{---}\kern-.5em--- x

Why not `\hbox{---\kern-.5em---}` or just `---\kern-.5em---` to get the same effect? This also assumes that the font collapses these three hyphens into a dash, then it backtracks the symbol width and does a second one.¹⁴ Anyway, where figure dashes are related to minuses we can probably ignore this super minus resembling horizontal bar.¹⁵

The Wikipedia page ends with a summary of all kind of dashes, including underscores, script specific symbols, accents (like macron), modifiers and curly ones. Here we only

¹³ The math stream doesn't go through the font handler although embedded `\hboxes` get that treatment. This means that two hyphens in a row are just two atoms and not get collapsed to an endash.

¹⁴ Here is some food for thought: for this kind of usage one can argue that such a dash should have some stretch. In LuaMetaTeX and therefore ConTeXt we can do this: `\uleaders \hbox to 1.5em {---\hskip Opt minus .5em---} \hskip.125em minus .125em \relax` and get: x — xx — xx xx — xx xx — x. Boxed material can be stretched and be taken into account when creating paragraphs. It is no big deal to wrap that in a macro, say `\figuredashed`.

¹⁵ We can actually issue a warning when it is used in math mode.

mention the ones that can end up in some source when one cuts and pastes. Doing that can result in missing characters (because not all fonts provides them) or a change in meaning (for as far as the symbols relates to an intention). We show some that fit into this discussion and also mention the Unicode description:

U+002D	HYPHEN-MINUS	the usual hyphen but also used as minus
U+005F	LOW LINE	aka underscore
U+00AD	SOFT HYPHEN	valid hyphenation point (invisible)
U+2010	HYPHEN	the real hyphen but more work on a keyboard
U+2011	NON-BREAKING HYPHEN	a hard hyphen, disables following hyphenation
U+2012	FIGURE DASH	see discussion above
U+2013	EN DASH	see discussion above
U+2014	EM DASH	see discussion above
U+2015	HORIZONTAL BAR	see discussion above
U+2043	HYPHEN BULLET	used in itemized lists
U+207B	SUPERSCRIPT MINUS	combined with pre-superscripted characters
U+208B	SUBSCRIPT MINUS	combined with pre-subscripted characters
U+2212	MINUS SIGN	the math minus (rendering of hyphen)
U+23AF	HORIZONTAL LINE EXTENSION	build long connected horizontal lines
U+23E4	STRAIGHTNESS	represents line straightness in technical context
U+2500	BOX DRAWINGS LIGHT HORIZONTAL	part of the box-drawing repertoire
U+2796	HEAVY MINUS SIGN	a visual variant with no meaning
U+2E3A	TWO-EM DASH	a visual variant with no meaning
U+2E3B	THREE-EM DASH	a visual variant with no meaning
U+FE58	SMALL EM DASH	a visual variant with no meaning
U+FE63	SMALL HYPHEN-MINUS	a visual variant with no meaning
U+FF0D	FULLWIDTH HYPHEN-MINUS	a visual variant with no meaning

The Unicode math draft only mentions the hyphen:¹⁶

“Minus sign. U+2212 [or] – [known as] MINUS SIGN is the preferred representation of the unary and binary minus sign rather than the ascii-derived U+002D [or] – [known as] HYPHEN-MINUS, because minus sign is unambiguous and because it is rendered with a more desirable length, usually longer than a hyphen.”

and elsewhere we can read:

“The ascii hyphen minus U+002D [or] – is a weakly mathematical character that may be used for the subtraction operator, but U+2212 [or] – [known as] MINUS SIGN is preferred for this purpose and looks better.”

¹⁶ When I copy this snippet into the document source there are START OF TEXT symbols at the places where a hyphenation occurs, which is probably a side effect of a bad TOUNICODE entry in the pdf file, but it is kind of interesting in this perspective as definitely a hyphen is rendered.

We are not aware of the concept of weak mathematical characters, so we will not take that property too serious when we try to improve the rendering.

This is basically it. There is no mentioning of classes (after all, traditional \TeX has no unary class) so it is assumed that the renderer does the right thing: interpreting the sequence of characters and apply spacing accordingly. There are users who like to see a unary minus being rendered differently, just as the minus that a student is supposed to key in a calculator and while the Wikipedia page mentions this explicitly, it is ignored here. Yes, having two distinctive slots for this would have been great. Maybe it is not seen as relevant enough by the community that would benefit most, but who knows what had happened if the Wikipedia page had been there before!

The minus is mentioned in the somewhat curious section about how shapes should be positioned relative to the baseline, where the position of the minus relates to what in \TeX speak is the math axis. There is also some mentioning of non-mathematical use, like:

*"The concept of mathematical use is deliberately kept broad; therefore the Math property is also given to characters that are used as operators, but are not part of standard mathematical notation, such as **U+2052 COMMERCIAL MINUS**."*

There should be no confusion with the **SET MINUS** which renders as a backslash, a **(NEG\-\ATED) MINUS TILDE** or **(NEG\-\ATED) SIMILAR MINUS SIMILAR** that look more like relations. **MPS: overfull hbox, and do you intend to hyphenate?**

The MathML document recognizes the minus as being unary or binary. In content MathML it is easy: when applied to a single atom it is a unary. In presentation MathML minus is an operator that sits at the front of a row (unary) or in the middle (binary). Keep in mind that we are limited to **mn** for numbers, **mi** for alphabetic symbols and **mo** for operators, not to be confused with \TeX 's math operators, because in MathML relations are also operators. One can wonder about a minus in **mn** elements.

So to summarize: we definitely need to make sure that (whatever renders as) hyphens is dealt with in math as minus. We can wonder what to do with (especially) en- and emdashes and the other horizontal lines that actually might show up as (what we call) middle delimiters in mathematical constructs: if it's there, \TeX ies will use it! The lack of specific symbols for unary minus has to be compensated at the macro package level.

16.7 Pieces

In Unicode one can find all kind of constructors, for instance characters that find their origin in those character sets that had lines and corners for drawing on a terminal. It is therefore no surprise that there are also some constructors that relate to math. An example demonstrates this:

```
\def\makeweird#1#2#3#4%
{\vcenter\bgroup
```

```

\offinterlineskip
\hbox{$\scriptscriptstyle\char"1$\par
\hbox{$\scriptscriptstyle\char"2$\par
\hbox{$\scriptscriptstyle\char"3$\par
\hbox{$\scriptscriptstyle\char"4$\%
\egroup}

\def\lwA{\mathopen {\makeweird{23A7}{23A8}{23A8}{23A9}}}
\def\rwA{\mathclose{\makeweird{23AB}{23AC}{23AC}{23AD}}}
\def\lwB{\mathopen {\makeweird{23A7}{23AC}{23AC}{23A9}}}
\def\rwB{\mathclose{\makeweird{23AB}{23A8}{23A8}{23AD}}}
\def\lwC{\mathopen {\makeweird{23A7}{23AC}{23A8}{23A9}}}
\def\rwC{\mathclose{\makeweird{23AB}{23A8}{23AC}{23AD}}}

$\lwA x + 4 + \lwB x^2 + 4^2 + \lwC x^3 + 4^3 \rwC \rwB \rwA$
```

This renders as:

$$\left\{ x + 4 + \left\{ x^2 + 4^2 + \left\{ x^3 + 4^3 \right\} \right\} \right\}$$

So, we have official Unicode characters for constructing large fences. In the Unicode math documents there is some mentioning of this and interesting is that there are suggested compositions expressed in 2, 3, 5 etc. stacked ‘*lines*’ which makes one wonder how math is perceived (or supposed to be rendered). But what is really weird is that there are plenty of arrows but no snippets defined that can be used to create extended ones. Why vertical snippets and no horizontal ones? This is clearly an omission and the \TeX community did take care of this need. So, for horizontal arrows and alike one expects the font to handle it and for fences not?

It is not only fences that have snippets, we also find them for integrals. But for vertical arrows they are lacking: that is completely up to the font. Now, for us that is fine, but again, for consistency they could have been there. It would make it possible to filter bits and pieces from fonts using official slots instead of private ones. So, to some extent we can best assume there is nothing like that and ignore whatever pieces are in Unicode anyway (like the braces in the example). One can even argue that because of this inconsistency a font designed can as well only use private slots and not provide snippets at all.

So, how do we get out of this situation? Because no one cared getting it in Unicode, we can do as we like. Of course, we can define arrow fillers as has always been done in \TeX , but because in $\text{LuaMeta}\text{\TeX}$ we have a bit more in our toolkit, and because we want to support stretch fractions (where the rule is replaced by a horizontal delimiter) it was decided to define a tweak that deals with this: when the basic arrows have no horizontal parts defined, we just assemble them. For those arrows that have a hook or so at

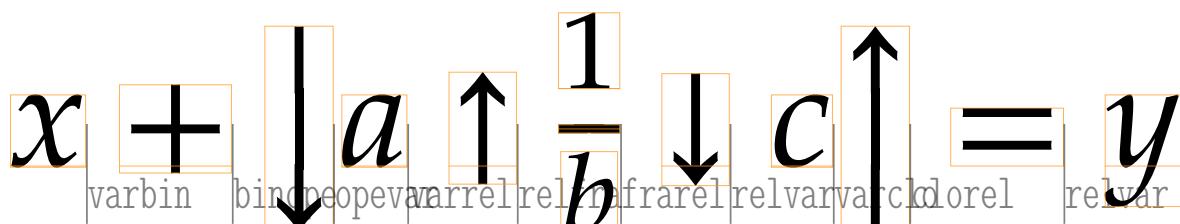
the other end, we use the space as extender.¹⁷ If we ever end up with proper snippets in Unicode then we also need adapted fonts, and then we can get rid of these hacks. That said: because all decent math fonts do have the three pairs of fences (brace, parenthesis, bracket) the vertical snippets are rather useless, unless one wants to construct assembled weird ones. This would be different for horizontal assemblies, because there is more variety in them.

The official name for all related to characters that can stretch is ‘*delimiter*’. In traditional TeX one can define a command that becomes a character. In that case a family, class and slot is assigned. You can also directly access a character in which case one will assign these properties otherwise (no command is defined). The same is true for these delimiters. However, in traditional TeX the larger character usually comes from a so called extension font and uses family 3). In OpenType fonts we have all in one font so there the large family, class, and slot are not used.

An interesting side effect of the updated math machinery in LuaMetaTeX is that we no longer really need delimiter specifications when we use OpenType fonts. This is because in practice the only two classes that really matter are the open and close ones. There are basically two kinds of delimiters: fences and singulars. Fences need open and close and only bars have a dual character. So, when we don’t define it as delimiter, the engine can still use that character and take its assigned class when used stand-alone, while in the case of fences these themselves are of class open and close. And, for instance a left brace can get class open because when used stand alone it is an unscaled left fence. In the rare case that one really need a different class we are using commands: some characters can be binary, ordinary or whatever so then commands relate a name to a class-character combination. Actually, in ConTeXt we will switch to using dictionaries and field specific rendering instead, but that is a different story. We can illustrate the arrows with an example:

```
$ x +
  \left\downarrow\right. a \uparrow \frac{1}{b} \left\downarrow\right. c \right\uparrow\left. y
= y $
```

The stand alone arrows are defines with class relation but when used as fences their spacing is driven by the fences themselves.



This means that in ConTeXt LMTX we no longer have delimiter code definitions. Of course the engine has to be able to use math characters of any kind (by commands,

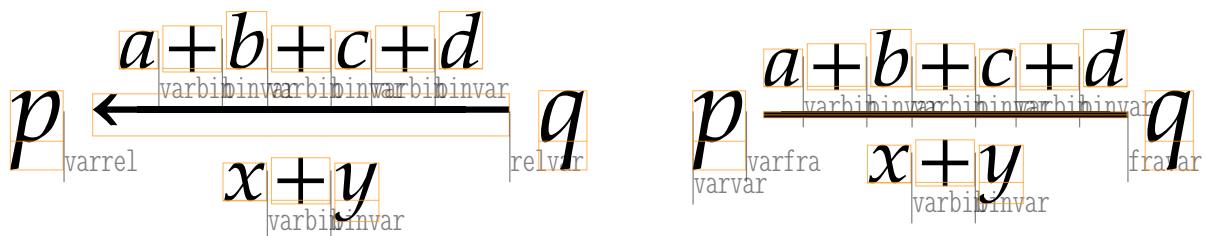
¹⁷ Actually we no longer do that because the engine will center the arrow anyway when it’s too short.

direct or as utf) as delimiters, but that was not that hard to provide. It also simplifies the code we use for fencing as it can be less selective.

Another interesting side effect of once again looking into these stretched characters is that the fraction mechanism that already was extended with skewed fractions, now supports any stretchable character as alternative for a fraction rule.

```
$
p \leftarrowtext {a + b + c + d}{x + y} q
\quad
p \frac {a + b + c + d}{x + y} q
$
```

Watch the difference in spacing: here the class of the used delimiter determines the spacing around the (pseudo) fraction:



Again this simplifies some code because normally one ends up with stacking stuff using leaders in between.

16.8 Accents

When we talk about accents, we refer to tiny symbols that anchor themselves onto base characters. We limit ourselves to the ones common in Latin scripts because they are the ones used in math. Accents in Unicode are somewhat special. In the past, when encoding vectors were limited, accents were entered as part of an input sequence and then anchored by the renderer. Nowadays often pre-composed characters are used. A very cheap way of anchoring is to have accents that just overlay, and in practice centering an accent over a base character works sort of okay. As an example of an accent we will use the hat:

U+005E	x^x	m^m	\Hat	$x^x + m^m$
U+02C6	\hat{x}^x	\hat{m}^m	\hat{}	$x^{\hat{x}} + m^{\hat{m}}$
U+0302	\hat{x}	\hat{m}	\widehat	$\hat{x}^x + \hat{m}^m$

Normally the font handler will take care of anchoring U+0302, but it can only be done properly when there are anchors defined for what are called ‘marks’: the official feature description is mark-to-base (or simply **mark**). The last column in the above table shows math and as we input a raw character we don’t get proper anchoring: the zero width makes it overlap.

Now wait, you will say, but why does it actually overlap? The reason is that zero width is not actually zero width here! The glyph has a bounding box that goes into the negative

horizontal direction and therefore, when such a shape gets injected into the output, the rendering in the viewer will move the left edge to the left. But because the \TeX engine only handles positive widths and because the width is explicitly part of a character specification anyway¹⁸ we don't progress (advance) which is why the overlapping sort of works for the x but less so for the m : in math mode we need to use these `\hat` and `\widehat` commands.

The hat and widehat assignments were those of August 2022. In plain \TeX we see these definitions:

```
\def\hat{\mathaccent"705E}
\def\widehat{\mathaccent"0362}
```

The `\mathaccent` primitive takes an integer that encodes the class, family, and slot in the 8 bit font encoding. Here we see that the hat comes from family 0, the upright math font. The widehat comes from extensible family 3. These two are independently defined. When you want a hat that spans the nucleus, you need to use the widehat. In the math engine spanning actually means that we have a delimiter and normally that means: start with a basic shape, when that is too narrow, go to the extensible font and follow the chain with increasing sizes and when you run out of those apply an extensible recipe. The sequence and extensible are both optional and the important part is that we first look at what is called the small character and then to the large one(s).

However, the `\mathaccent` primitives doesn't take a delimiter! It directly starts following a chain if the given character has it (and then the character itself is of course the first in that chain). And this is where the problems start when we move to OpenType and Unicode math.

U+005E	Hat	some useless, often ugly large glyph
U+02C6	hat	it has width but no extensibles
U+0302	widehat	it has zero width and extensibles

Now, if we define `\hat` as U+02C6 we don't get the extensibles, and it basically is what was always done in \TeX macro packages following the plain suggestions. If we define `\widehat` we start out with a glyph that has likely zero width¹⁹ And, because OpenType starts with the base glyph and *then* uses a set of variants of eventually a recipe of parts, we suddenly have a different situation with `\mathaccent` than we normally have, where these are decoupled. Therefore, the definition of `\hat` and `\widehat` determines what an OpenType math engine will do, just as in regular \TeX , but we might need them to be defined differently.

A solution would be to let `\mathaccent` (or `\Umathaccent`) directly go to the variants, but that is sort of weird. Because a zero width glyph doesn't match the criteria to span a nucleus it is likely to be skipped anyway, although there can be a case where the next in

¹⁸ The height and depth are not: these we derive from the bounding box.

¹⁹ Over the many years that \TeX evolved this was not guaranteed, for instance when wide (Unicode) fonts were constructed from traditional eight bit (\TeX encoded) fonts.

size overruns the width of the nucleus in which case the zero width one is used which itself is not that nice. We could actually derive the width from the boundingbox, but that would be a bit abnormal, and it makes no sense to burden the font machinery with that exception. Another approach we can follow is to just copy the extensibles from `U+0302` to `02C6` and use that one for `\hat` as well as `\widehat` and then make `\widehat` an alias to `\hat`. After, all, the main reason why we have two commands comes from the fact that `\mathaccent` doesn't take a delimiter but single character reference (encoded in an integer).

Here is the whole list of accents:

<code>\grave</code>	<code>U+0060</code>	<code>\widegrave</code>	<code>U+0300</code>
<code>\ddot</code>	<code>U+00A8</code>	<code>\widedot</code>	<code>U+0308</code>
<code>\bar</code>	<code>U+00AF</code>	<code>\widebar</code>	<code>U+0304</code>
<code>\acute</code>	<code>U+00B4</code>	<code>\wideacute</code>	<code>U+0301</code>
<code>\hat</code>	<code>U+02C6</code>	<code>\widehat</code>	<code>U+0302</code>
<code>\check</code>	<code>U+02C7</code>	<code>\widecheck</code>	<code>U+030C</code>
<code>\breve</code>	<code>U+02D8</code>	<code>\widebreve</code>	<code>U+0306</code>
<code>\dot</code>	<code>U+02D9</code>	<code>\widedot</code>	<code>U+0307</code>
<code>\ring</code>	<code>U+02DA</code>	<code>\widering</code>	<code>U+030A</code>
<code>\tilde</code>	<code>U+02DC</code>	<code>\widetilde</code>	<code>U+0303</code>
<code>\ddot</code>	<code>U+20DB</code>	<code>\widedot</code>	<code>U+20DB</code>

The only accent that is an exception is the last one but is it really used? It anyway makes no real sense to assume that users will ever directly input the utf characters conforming the last column, so we can just go for the first one and use the extensibles from the second and see where we end up. Neither MathML nor TeX related specifications seem to cover this well, so we can just do what suits us best.

Because all has to fit into the ConTeXt user interface and because we also want to be backward compatible (command wise), we end up with something:

```
\showglyphs
\im {\widehat{a} + \hat{a} {aa}} =
\im {\hat{a} {aa} + \hat{a} {aa}} =
\im {\hat{a} {aa} + \hat{a} [stretch=yes] {aa}} =
\setupmathaccent [top] [stretch=yes]
\im {\hat{a} {aa} + \hat{a} {aa}}
```

that gives us:

$$\widehat{a} + \widehat{aa} = \widehat{a} + \widehat{aa} = \widehat{a} + \widehat{aa} = \widehat{a} + \widehat{aa}$$

Now, one problem, is of course that users can enter these modifiers as utf sequence in the input, just like they do with delimiters. Therefore we do support the following feature (which is under class control so disabled by default):

```
\Umathcode "02C6 \mathaccentcode 0 "02C6
```

```
\edef \HiHatA {\Uchar"02C6}
\Umathchardef \HiHatB \mathaccentcode 0 "02C6

$ \Uchar"02C6{x} + \HiHatA{xx} + \HiHatB{xx} = \widehat{xxxx} $
```

You get this:

$\widehat{xx} + \widehat{xx} = \widehat{xxxx}$

The only cheat here is that normally accents come after the accentee, but we can live with that. After all, it's all about convenience.

There is another aspect of accents that we need to mention here. The hat, tilde, and check are often used over not only single letters but also small expressions. So how come that fonts have only very few variants defined? We can imagine that in eight bit fonts the number of available slots plays a role but in OpenType fonts that is not the case. It therefore can be considered an oversight that usage of these wide accents has not been communicated well to the font designers.

$$\begin{aligned} \hat{a} &+ \overbrace{a+b} + \overbrace{a+b+c} + \overbrace{a+b+c+d} + \overbrace{a+b+c+d+e} + \overbrace{a+b+c+d+e+f} \\ \tilde{a} &+ \overbrace{a+b} + \overbrace{a+b+c} + \overbrace{a+b+c+d} + \overbrace{a+b+c+d+e} + \overbrace{a+b+c+d+e+f} \\ \check{a} &+ \overbrace{a+b} + \overbrace{a+b+c} + \overbrace{a+b+c+d} + \overbrace{a+b+c+d+e} + \overbrace{a+b+c+d+e+f} \end{aligned}$$

The previous lines demonstrate that we can actually cheat a little for these three top accents: we can just scale the last variant horizontally. It was a few lines patch to Lua-MetaTeX to make this automatic and triggered by setting the `extensible` field in a character table to `true` instead of a recipe. The ingredients to get this working were already there, and it works out quite well. The only complication was that the `flac` feature (that provides flat accents for cases where the nucleus is rather high) could interfere, but that was trivial to deal with in the code that does the goodies.²⁰

When it comes to these delimiters that have no real solution in the font, we can consider delegating coming up with a glyph to the macro package at the time it is needed, and we can actually do that. However, this is mostly interesting for educational usage, where the amount of delimiters is predictable and limited. About a decade ago some mechanism was added to the MkIV math machinery that support plugins so that we could use MetaFun to generate (most noticeably) square root symbols the way we liked.²¹ The main drawback is that mixing this in means matching to a font, and that is not always trivial. But it is this kind of trickery that makes working with TeX fun. That said: what we are discussing here is more fundamental in the sense that we try to come up with generic engine solutions that just rely on the fonts. That way complex math with all reasonable symbols is also served.²²

²⁰ When we were testing fonts this got us by surprise when we tested Cambria that has these flat overloads for the tilde and check. Because supports this automatic (hidden from the user) one doesn't look into that direction when testing something.

²¹ This was a fun project of Alan and Hans.

²² These MetaFun plugins are still possible, but we need to adapt some to LMTX which will happen as we go.

Interestingly there are some arrows that act like accents. There are over- and under ones as well as combining (often zero width) accents. Fonts are not always consistent in how these extends (the wide ones). Often the combining accents are smaller and closer to the running text. Traditionally in \TeX fonts there are no extensible arrows: they are constructed from arrow heads, minus and equal signs with some negative spacing in between. One can therefore wonder is the smaller combining ones are appreciated by those who want stable math. It definitely means that we have to make choices. Even more interesting is that while Unicode has some means to construct braces from predictable Unicode slots. there is no way to do the same with arrows and (indeed) there are fonts out there with shaped arrows that demand different middle and end pieces. In fact, the same is true for rules that are not simple rectangles and radical extensions that are not flat rules either. In all these cases the usage patterns of accents and similar constructs has not really been fed back into the way Unicode and OpenType fonts support math.²³

16.9 Bullets

In \TeX usage bullets are a bit special. Because fonts had a limited number of slots available, bullets in for instance itemized lists traditionally were taken from a math font. The bullet in Computer Modern has a comfortable size and is quite useful for that. Bullets in text fonts often were (are) relatively small so even when they were available they were not really used. The official Unicode slot for bullet is **U+2022** and in this font it shows up as ‘•’. The Wikipedia page on bullets (typography) mentions:

*“A variant, the bullet operator (**U+2219 • BULLET OPERATOR**) is used as a math symbol, akin to the dot operator. Specifically, in logic, $x \bullet y$ means logical conjunction. It is the same as saying “ x and y ””*

The page also mentions that “*glyphs such as • and ○*” have “*reversed variants □ and ▨*” although we haven’t seen the reverse once in \TeX documents (yet), like these (we use Stix2 to show them):

U+2022	•	BULLET
U+2023	?	TRIANGULAR BULLET
U+2043	-	HYPHEN BULLET
U+204C	«	LACK LEFTWARDS BULLET
U+204D	»	LACK RIGHTWARDS BULLET
U+2219	•	BULLET OPERATOR (math)
U+25CB	○	WHITE CIRCLE
U+25CF	●	BLACK CIRCLE
U+25D8	▣	INVERSE BULLET
U+25E6	○	WHITE BULLET

²³ One can argue that this is not what Unicode is for but if so, then some other bits and pieces also make little sense.

U+29BE ◎ CIRCLED WHITE BULLET

U+29BF ◉ CIRCLED BULLET

The reverse ones are not really reverse in Stix2 as they have bigger circles. There are a few more bullets mentioned but probably only because they have the word bullet in their description and they don't really look like bullets. Given the already discussed lack of granularity in some math symbols with multiple usage it is somewhat surprising that we have a math bullet. The weird looking left- and rightward bullets are kind of hard to distinguish. Let's hope that mathematicians don't discover these!

This brings us to the more general way of looking at these bullets because among the popular math symbols used in text are also the triangles and (T_EX) math fonts came with. When we have a few commands for circular shapes like `\bullet\big-circ\circ` giving • ○ ◦ we have plenty of (black) triangles.

For instance, we have `\triangledown` and `\bigtriangledown` and these have corresponding Unicode slots U+25BD and U+25BF but when you try these in for instance Stix2, Pagella and Cambria you got: ▽ + , ▽ + ? and ? + ?, where the question mark indicates a missing character.

It is for that reason that `\triangledown` and `\bigtriangledown` are both defined as using the large one. This test also demonstrated us that we didn't have to waste time looking up what MathML had to tell about it. A typeset version of that specification was never a visual highlight and missing glyphs only makes that worse. And, when fonts lack shapes no one uses them anyway.

However, it makes sense to think a bit about how to deal with this properly, and we will likely add some checking to the goodie files for it, so that when we do have them, we use them.²⁴ But even then, most troublesome is that the size (and even positioning) of these symbols is rather inconsistent across math fonts, but because they are seldom used it doesn't make much sense to compensate for that (read: we just wait till users ask for it).

16.10 Punctuation

There are quite some punctuation symbols in Unicode but not for math where the main troublemakers are the period, comma, colon and semicolon. The first two can be used as separator in numbers, in which case we don't want any spacing, or they can be part of a (pseudo) sentence in a formula, or they can separate entries in a list (take coordinates).

1.1 + 1.2

(1.1, 1.2)

x + 1.1, x + 1.2

When used as separator in a sentence, which is more likely in display math than in inline math, the spacing after it can be either regular (as in text) or wide. And the

²⁴ Most practical is to add this information to the character database which is a bit of work

symbol can come from the math font or text (and these can actually look different). In ConTeXt (also pre LMTX) we have some special trickery at work for spacing comma's and periods but we leave that aside now. What should be noted is that out-of-the-box spaces are ignored when math is scanned so we cannot take that surrounding into account when dealing with spacing in the engine.

Although the Unicode specification provides a classification of characters that includes punctuation in practice we need to deal with it ourselves. For instance, by default a period is not considered punctuation but a command and semi colon are, while a colon is a relation!

Take for instance *f.* (math italic f followed by a period). Italic correction and math glyphs have this special relationship and it also shows up in punctuation. Imagine that we have a sequence of characters, say *fx*. These are actually two ordinary atoms but in *f.* we have an ordinary atom followed by a punctuation atom so here spacing is determined by how these classes are set up. But, given the shape of the *f* we actually don't want italic correction here.

$fx + f. + f, + f : + f; + a. + a, + a : + a; + x, + x, + x : + x;$

When you zoom in you can see the subtle spacing differences. We can compensate for the semi colon being a bit higher than the period by applying some kern, something that we can set up in the goodie file.

Actually, if we assume that periods only occur in numbers we can make it punctuation and set it up for digit spacing but then commas etc also get done that way. A variant is to have two punctuation classes (or cheat and put the period in the digit class). No matter what we do, no help can be expected from documents mentioned: it's mostly a visual thing anyway.

Let's end with the visual aspect: in most fonts the two colons `0x003A` and `0x2236` are different: one has more distance between the periods. Which one? Well, that depends on the font! Latin Modern has a cramped `0x2236` while Stix2 has a cramped `0x003A`. Cambria has square dots for the `0x003A` and round ones slightly more cramped for `0x2236`. Lucida goes extreme: it has smaller dots far apart for `0x2236`. If the idea is that a reader should get from the shape what it's about one can wonder if texts get read the way the author intended. Of maybe shapes don't matter. Of course a macro package can obscure these inconsistencies by setting the math character code of `0x003A` to `0x2236` but that only obscures the fact that little attention has been paid: what one can consider bugs became features.

16.11 Special ones

There are quite some characters that really depend on a math renderer. Examples are wide accents, fences, and arrows. Some constructs, like fractions use rules and these

don't come from Unicode nor fonts. A mixed case is radicals: there is a Unicode point and fonts can provide larger variants. Normally one steps up a slightly slanted version but when things get large the radical becomes an extensible and therefore gets an upright shape. The engine is supposed to add a horizontal rule at the right location. Interesting is that there is no provision for a right end cap. The reason probably is that \TeX , being the major renderer, has no combined horizontal and vertical extenders and OpenType doesn't have that either. Some properties are driven by the fonts' math parameters which sort of makes the radical rendering a very restricted adventure: it is supposed to be used for roots only, either or not with a degree anchored in the right top area. It looks like that degree is not really to extend much beyond the left edge of the symbol.

In Unicode there is an actuarian character **U+20E7** and support in fonts is not that good. We do support it because we ran into it in MathML. However, it is a hack. The symbol as provided by fonts is rather useless.

```
$ \sqrt{x + 1} + \annuity{x + 1} $
```

Let's see how it renders:

$$\sqrt{x + 1} + \overline{x + 1}$$

We take the dimensions of a radical as template and when we look at the bare glyphs we see this:

$$\sqrt{} \quad \overline{}$$

Basically we have a right actuarian character like we have a left radical. But In this case the rule will go left instead of right. This is implemented on top of radicals so and driven by **\Udelimited** that takes two delimiters and doesn't scan for a degree. For two-sided roots (with degree) we have **\Urooted**. And like normal radicals the delimited one adapts itself to the content:

```
$ \sqrt{x + \frac{1}{x}} + \annuity{x + \frac{1}{x}} $
```

So we get:

$$\sqrt{x + \frac{1}{x}} + \overline{x + \frac{1}{x}}$$

For the record: in ConTeXt spacing is also driven by the struts and because we use the radicals renderer the gap and distance parameters also apply. It might look spacy, but keep in mind that we want radicals to look similar when we have more of them in line,

and we can configure all. We have also enabled the feature that radicals at the same level are normalized in height and depth. Here are some variants:

```
$ \lannuity {x + \frac{1}{x}} +  
 \rannuity {x + \frac{1}{x}} +  
 \lranannuity {x + \frac{1}{x}} $
```

This gives:

$$\overline{x + \frac{1}{x}} + \overline{x + \frac{1}{x}} + \overline{x + \frac{1}{x}}$$

So we can have a mix of left, right and both end radical like symbols that encompass the nucleus. We're not aware of more such characters in Unicode but when they show up we are prepared. Only real usage can result in some parameters being fine-tuned.

16.12 Final words

This text was written in 2022 when we were working on math, extending the goodie files with new tweaks, checking support in fonts and updating manuals. But, as we moved forward, for instance with adapting Type1 support of Antykwa and Iwona to the new possibilities again we had to go back in time and figure out why actually things were done in certain ways. And I have to admit that we had some good laughs and quite some fun on seeing how strange and inconsistent the assumed structured and logical \TeX ecosystem deals with math. A wrapup like is is never complete and we can keep adding to it so just consider it to be a momentary impression.

Personally I have to admit that I've always overestimated what happened outside the $\text{Con\TeX{}t}$ bubble, especially given the claims made. Consistency in Unicode math is probably not as good as is could have been and the same is true for OpenType math support, but maybe I'm naive in expecting consistency and logic in math related work. The mere fact that Donald Knuth pays a lot of attention to the math in his writing doesn't automatically translate in all $\text{\TeX}ies$ doing the same. I don't claim that $\text{Con\TeX{}t}$ is doing better but I do hope that its users keep going for the best outcome.

One final note. In $\text{Con\TeX{}t}$ we always tried to keep up with developments and Unicode input as well as using OpenType math fonts are part of that. However, because we're not part of the '*gremia of \TeX math and related coding*' it hardly matters what our opinions are with respect to these issues. The best we can do is adapt to whatever shows up, it being bad or good. It is however kind of funny to see (by now rusty) problems that have been noticed already long ago being presented as kind of new. Hopefully staying ahead and/or adapting with specific solutions doesn't backfire to hard on the $\text{Con\TeX{}t}$ users. If so, we're sorry for that. As long as they can render their documents well, it doesn't matter that much anyway. After all, we can always just blame '*the others involved*'.

16.13 Resources

- [1] [en.wikipedia.org/wiki/Slash_\(punctuation\)](https://en.wikipedia.org/wiki/Slash_(punctuation))
- [2] www.unicode.org/reports/tr25
- [3] www.w3.org/TR/MathML3
- [4] www.unicode.org/Public/math/revision-15/MathClass-15.txt
- [5] en.wikipedia.org/wiki/Vertical_bar
- [6] en.wikipedia.org/wiki/Dash
- [7] en.wikipedia.org/wiki/Commercial_minus_sign
- [8] en.wikipedia.org/wiki/Division_sign
- [9] [en.wikipedia.org/wiki/Bullet_\(typography\)](https://en.wikipedia.org/wiki/Bullet_(typography))

Starting with TeXLive 2023 the default ConTeXt distribution is LMTX, a follow up on MkIV, running on top of the LuaMetaTeX engine instead of LuaTeX. Already for a long time the MkII version used with pdfTeX, XeTeX and Aleph has been frozen and most users moved on from MkIV to LMTX (a more distinctive tag for what internally is version MkXL).

In principle one can argue that we now have three versions of ConTeXt and there can be the impression that they are very different. However, although MkXL can do more than MkIV which can do more than MkII, the user interface hasn't changed that much and old functionality is available in newer versions. Of course some old features make no sense in newer variants, like eight-bit font encodings in an OpenType font realm and input encodings when one uses utf, although we still support input encodings a.k.a. regimes. When we started using the **Mk*** suffixes the main reason was that we had to distinguish files and the official TeX distribution doesn't permit duplicate file names. Using a distinctive suffix also makes it possible to treat files differently.

suffix	engine	template	arguments	main file
MkII	pdfTeX, XeTeX, Aleph			context.mkii
MkIV	LuaTeX, LuajitTeX, LuaMetaTeX			context.mkiv
MkVI	idem		yes	
MkIX	idem	yes		
MkXI	idem	yes	yes	
MkXL	LuaMetaTeX			context.mkxl
MkLX	idem		yes	

In this table '*template*' files are a mix of TeX and Lua and originate in the early days of MkIV; basically, they are a wink to active server pages. With '*arguments*' we refer to files that accept named macro arguments which means that they need to be preprocessed. That started as a proof of concept but some core files are defined that way. Users will normally just use a **.tex** file.

The Lua files in the code base have the suffix **lua**, or when meant for LuaMetaTeX that uses a newer Lua engine they can have the suffix **lmt**. There can also be **lfg** (font goodies) and **lg** (language goodies) plus byte-compiled files with various suffixes but these are normally not seen by users. We leave it at that.

So, while TeXLive 2022 installed MkII and MkIV, TeXLive 2023 installs MkIV and LMTX. Therefore the most significant upgrade is in the engine that is used by default: LuaMetaTeX instead of LuaTeX. The MkII files are no longer installed so we don't need pdfTeX.

So how did we end up here? Initially the idea was that, because LuaTeX is basically frozen, LuaMetaTeX would be the engine that we conduct experiments with and from

which occasionally we could backport code to LuaTeX. However it soon became clear that this would not work out well so backporting is off the table now. Just for the record: the project started years ago so we're not talking about something experimental here. There have been articles in TugBoat about what we've been doing over the years.

One of the first decisions I made when starting with LuaMetaTeX was to remove the built-in backend, which then meant also removing the bitmap image inclusion code. That made us get rid of dependencies on external libraries. In fact, a proof-of-concept experimental variant didn't use the built-in backend at all. The font loading code could be removed as well because that was not used in MkIV either. In MkIV we also don't use the kpse library for managing files so that code could be dropped from the engine tool; it can be loaded as so-called optional library if needed but I'll not discuss that here. If you look at what happens with the LuaTeX code base, you'll notice that updating libraries happens frequently and that is not a burden that we want to impose on users, especially because it also can involve updating build-related files. Another advantage of not using them is that the code base remains small.

A direct consequence of all this was that the build process became much more efficient and less complex. A fast compilation (seconds instead of minutes) meant that more drastic experiments became possible, like most recently an upgrade of the math subsystem. All this, combined with an overhaul of the code base, both the TeX and MetaPost part, meant that backporting was no longer reasonable. Being freed from the constraint that other macro packages might use LuaMetaTeX in turn resulted in more drastic experiments and adding features that had been on our wish list for decades. Another side effect was that we could easily compile native MS Windows binaries and immediately support transitions to arm-based hardware.

Instead of "*backporting after experimenting*", a leading motive became "*fundamentally move forward*" while at the same time tightening the relation between ConTeXt and the engine: the engine code became part of the distribution so that users can compile themselves, which fits perfectly in the paradigm (and demands) of distributing all the source code, even that of the engine. There is also less danger that patches on behalf of other usage interferes with stable support for ConTeXt. A specific installation is now more or less long-term stable by design because it no longer depends on binaries and/or libraries being provided for a specific platform and operating system version. Of course installers and TeXLive do provide the binaries, so users aren't forced to worry about it, but they can move along with a system update by recompiling an old, and for their purpose, frozen ConTeXt code base.

An unofficial objective (or challenge) became that the accumulated source stays around 12 MB uncompressed, (compressed a bit over 2 MB) and the binary around 3 MB so that we could use the engine as an efficient Lua runner as well as a launcher stub, thereby removing yet another dependency. That way the official ConTeXt distribution didn't grow much in size. A bonus is that we now use the same setup for all operating systems. It also opened up the possibility of a exceptionally small installation with all bells and whistles included. Another nice side effect, combined with automatic compilation on the compile farm, makes that we can provide installations that reflect the latest state

of affairs: a recent binary combined with the latest ConTeXt. As a result, most users quickly went for LMTX instead of MkIV.

In the code base we avoid dependencies on specific platforms but there are a few cases where the code for MS Windows and Unix differs. However, the functionality should be the same. A good test is that for MS Windows we can compile with mingw (cross-compilation), MSVC (native) and clang (native); that order is also the order of runtime performance. The native MSVC binary is the smallest but users probably don't care. In any case, it is nice to have a fallback plan in place. The code is all in C; the MetaPost code is converted from cweb into C using a Lua script but we also ship the resulting C code. The code base provides a couple of cmake files and comes with a trivial build script.

When I say that there are no libraries used, I mean external libraries. We do use code from elsewhere: adapted `avl` as well as `decnrumber` (for the MetaPost library), adapted `hjn` (hyphenation), `miniz` (zip compression), `pplib` (for loading pdf files), `libcerf` (to complement other math library support, but it might be dropped), and `mimalloc` for memory management. However all the code is in the LuaMetaTeX code base and only updated after checking what changed. The most important library originating elsewhere is of course Lua: we use the latest and greatest (currently) 5.4 release. We kept the `socket` library but it might be dropped or replaced at some point. In addition there is a subsystem for dynamically loading libraries; the main reason for that being that I needed `zint` for barcodes, interfaces to sql databases, a bunch of compression libraries, etc. But all that is tagged *optional* and ConTeXt will never depend on it. There are no consequences for compilation either because we don't need the header files. The glue code is very minimalistic and most work gets delegated to Lua.

Initially, because the backend is written in Lua, there was a drop in performance of some 15% but that was stepwise compensated by gains in performance in the engine and additional or improved functionality. The ConTeXt code base is rather optimized so there was little to gain there, apart from using new features. Existing primitive support could also be done a bit more efficiently; it helps if one knows where potential bottlenecks are. Therefore, in the meantime an LMTX run can be quite a bit faster than a MkIV run and it can even outperform a LuajitTeX run. In practice, the difference between an eight-bit MkII run using the eight-bit pdfTeX engine and a 32-bit LuaMetaTeX run with LMTX can be neglected, definitely on more complex documents. I never get complaints about performance from ConTeXt users, so it might be a minor concern.

So what are the main differences in the installation? If you really want to experience it you should use the standard installation. Currently the small installer is the engine that synchronizes the installation over the net and, assuming a reasonable internet connection, that takes little time. The installation is relatively small, and many of the bytes used are for the documentation. Updates are done by transferring only the changed files. The TeXLive installation is a bit larger because it shares for instance fonts with the main installation and these come with resources used by other macro packages. Both installations bring MkIV as well as LMTX and therefore provide LuaTeX as well as Lua-MetaTeX. However, a MkIV run is now managed by LuaMetaTeX because we use that

engine for the runner. The MkII code is no longer in $\text{\TeX}{}^{\text{Live}}$ but is in the repositories and used to test and compare with pdf $\text{\TeX}{}^{\text{}}^{\text{}}$. It just works.

The number of binaries and stubs is reduced to a minimum:

file	symlink	
tex/texmf-platform/luametatex		combined TeX, MetaPost and Lua engine
tex/texmf-platform/mtxrun	luametatex	script runner, binary
tex/texmf-platform/context	luametatex	CONTEXT runner, binary
tex/texmf-platform/mtxrun.lua		script runner, lua code
tex/texmf-platform/context.lua		loader for ConTeXt runner
tex/texmf-platform/luatex		the good old ancestor

All of these programs are in the ConTeXt distribution directory `tex/texmf-<platform>/`. In addition, `context` and `mtxrun` are symlinks to the `luametatex` binary, where possible.

So, the `context` command runs `luametatex`, but loads the Lua file with the same name which in turn will locate the ConTeXt management script (`mtx-context`) in the `TeX` tree and run it. The same is true for `mtxrun`: it is a binary (link) that loads the script in (this time) the same path and then can perform numerous tasks. For instance, identifying the installed fonts so that they can be accessed by name is done with:

```
mtxrun --script font --reload
```

Where in MkII we had stubs for various utility scripts, already in MkIV we went for a generic runner and a bit more keying. It's not like these scripts are used a lot and by avoiding shortcuts there is also little danger for a mixup with the ever-growing list of other scripts in TeXLive or commands that the operating system provides.

The LuaTeX binary is optional and only needed if a user also wants to process MkIV files. There are no shell scripts used for launching. The two main calls used by users are:

```
context foo.tex  
context --luatex foo.tex
```

A user has only to make sure that the binaries are in the path specification. When you run from an editor, the next command does the work:

```
mtxrun --autogenerate --script context <filename>
```

with `<filename>` being an editor-specific placeholder. Like other engines, LuaMeta- \TeX (and Con \TeX t) needs a file database and format file, and although it should generate these automatically you can make them with:

```
mtxrun --generate  
context --make
```

The rest of the installation is similar to what we always had and is tds compliant. The source code of LuaMetaTeX is included in the distribution itself (which nicely fulfills the requirements) but can also be found at:

<https://github.com/contextgarden/luametatex>

There are also some optional libraries there but ConTeXt works fine without them. The official latest distribution of ConTeXt itself is:

<https://github.com/contextgarden/context>

<https://github.com/contextgarden/context-distribution-fonts>

We see users grab fonts from the Internet and play with them. They can install additional fonts in `tex/texmf-fonts/data/<vendor>`. Project-specific files can be collected in `tex/texmf-project/tex/context/user/<project>`. These directories are not touched by installations and can easily be copied or shared between different installations. After adding files to the tree `mtxrun --generate` will update the file database.

In the distribution there are plenty of documents that describe how LuaMetaTeX with LMTX differs from MkIV with LuaTeX: new primitives, macro extensions, more granular math rendering, improved memory management, new (or extended) (rendering) concepts, more MetaPost features; most is covered in one way or another, and much is already applied in the ConTeXt source code. After all, it took a few years before we arrived here so you can expect substantial refactoring of the engine as well as the code base, and therefore eventually there is (and will be) more than in MkIV.

When you compare a ConTeXt installation with what is needed for other macro packages you will notice a few differences. One concerns the way TeX is launched. An engine starts with a blank slate but can be populated with a so-called format file that is basically a memory dump of a preloaded macro package. So, the original way to process a file is to pass a format filename to the engine. In order to avoid that a trick is used: when an engine (or symlink/stub to it) is launched by its format name, the loading happens automatically. So, for instance `pdflatex` is actually an equivalent for starting pdfTeX with the format file `pdflatex.fmt` while `latex` is pdfTeX with another format file (`latex.fmt`) starting up in dvi mode. And, as there are many engines, a specific macro package can have many such combinations of its name and engine.

In ConTeXt we don't do it that way. One reason is that we never distinguished between backends: MkII uses an abstract backend layer and load driver files at runtime (it was one of the reasons why we could support Acrobat as soon as it showed up, because we already supported the now obsolete but quite nice dviwindo viewer). And that model hasn't changed much as we moved on. Because we use a runner, we also don't need to distinguish between engines: all formats have the same name but sit on an engine subpath in the TeX tree. Anyway, this already removes quite some formats. On the other hand, ConTeXt can be run with different language specific user interfaces which means that instead of just `context.fmt` we have `cont-en.fmt` and possibly

more, like `cont-nl.fmt`. So that can increase the number again but by default only the English interface is installed. As a side note: where with MkII we needed to generate MetaPost mem files, with its descendants having `mplib` we load the (actually quite a bit of) MetaPost code at runtime.²⁵

In addition to a format file, for the `LuaTeX` and `LuaMetaTeX` engine we also have a (small) Lua loader alongside the format file. All this is handled by the runner, also because we provide extensive command line features, and therefore of no concern to users and package maintainers. However, it does make integrating `ConTeXt` in for instance `TeXLive` different from other macro packages and thereby puts an extra burden on the `TeXLive` team. Here I want to thank the team for making it possible to move forward this way, in spite of this rather different approach. Hopefully a `LuaMetaTeX` integration is a bit easier in the long run because we no longer have different stubs per platform and at least the binary part now has no dependencies and only has a handful of files.

For those new to `ConTeXt` or those who want to try it in `TeXLive` 2023 there is not much difference between the versions. However, MkIV is now frozen and new functionality only gets added to LMTX. Of course we could backport some but with most users already having moved on, it makes no sense. Just as we keep MkII around for testing with `pdfTeX`, we also keep MkIV alive for testing with `LuaTeX`. Maybe in a couple of years MkIV will go the same route as MkII: ending up in the archives as an optional installation.²⁶

²⁵ Occasionally I do experiments with loading the `TeX` format code at runtime, but at this moment the difference in startup time of about one second (assuming files are cached) is too large and running over networks will be less fun, so the format file will stay. The time involved in loading MetaPost can be brought down but for now I leave it as it is.

²⁶ This text appeared in TugBoat around the 2023 `TeXLive` release. Thanks to Karl Berry for his careful reading and fixing of the text and of course for keeping `TeXLive` alive.

How not to install ConTeXt 18

Installing LuaMetaTeX can be a complex process that requires some technical expertise, but the following steps should give you a general idea of what is involved:

1. First, you need to ensure that you have a recent version of the Lua programming language installed on your system. You can download the latest version of Lua from the official website at <http://www.lua.org/download.html>.²⁷
2. Next, you need to download the latest version of the MetaTeX distribution, which includes the LuaMetaTeX engine, from the official ConTeXt Garden website at https://wiki.contextgarden.net/ConTeXt_Standalone.²⁸
3. Once you have downloaded the MetaTeX distribution, extract the files to a directory on your system.²⁹
4. You can then run the LuaMetaTeX engine by opening a command prompt or terminal window and navigating to the directory where you extracted the MetaTeX files. From there, you can run the command "luametatex" followed by the name of the TeX file you want to process.³⁰
5. To make it easier to use LuaMetaTeX with your favorite text editor, you may also want to install a TeX distribution such as TeX Live or MiKTeX, which includes support for LuaMetaTeX. These distributions typically include a graphical user interface that makes it easier to manage your TeX installation and configure your system for use with LuaMetaTeX.³¹

It's worth noting that the exact steps for installing LuaMetaTeX may vary depending on your operating system and the specific TeX distribution you are using. For more detailed instructions, you may want to consult the official documentation for LuaMetaTeX and the TeX distribution you are using.³²

by ChatGPT^{33 34}

²⁷ The Lua code needed is part of the source tree that can be downloaded from GitHub or websites.

²⁸ There is no MetaTeX, although we sometimes joke about it.

²⁹ So here one is stuck.

³⁰ Shouldn't it be compiled first? And even then it needs some format, so one needs `context` and `mtxrun`.

³¹ Indeed installing TeXLive is easier, as is installing the smaller reference installation which uses LuaMetaTeX as its own installer. As far as we know, MikTeX doesn't include LMTX. And yes, consulting the documentation might be best.

³² So here is the way out of the proposed mess.

³³ Queried by Mikael Sundqvist.

³⁴ Which (at least here) is more about well formed sentences than about verified content. We can only hope that new TeX users are able to recognize a fake.

In the process of extending \TeX into LuaMetaTeX I occasionally run into what I consider side effects. Some might argue that we're seeing a bug but personally I never consider what I observe as "*interestingly weird*" to be a bug. After all, if it really was an issue Don Knuth would already have run into it and dealt with it accordingly.

Mid-2023 Mikael Sundqvist and I were reviewing the upgraded multi-line display math mechanism and in particular we were looking at how it behaves with for instance side floats. At some point we wondered if such a formula should follow the shape. This is not that trivial because where regular text is in most cases sort of grid-based and therefore has well-defined line heights, display math is way less predictable. This means that setting up hanging indentation based on lines will often lead to bad results. I did actually check if the `par builder` could be adapted to work with a height specification but this is not trivial, if only because one then needs to keep track of (intermediate) line heights while at the same time one doesn't know if the interline spacing changes, which actually is something a user can set up in ConTeXt display math. Setting the `hang` or `shape` comes before the content so we also have an out-of-order situation to deal with. Of course placing a formula number in the margin or below a formula also interferes.

Because formulas following a shape looked weird anyway (think for a moment what mid-aligned display math should look like) in the end we decided that a display formula in a side float situation is just an unbroken vertical box that if needed will extend below the allotted `hang` space. This is not that different from what \TeX does when we would use display math (which we don't in ConTeXt , where we follow a different approach) because \TeX just assumes that a display formula counts as three lines progression in a hanging (or `par` shape) continuation.

```
\dontleavehmode
\hbox{\strut one}
\break
\hbox{\strut two}
\par
\the\prevgraf
```

This example shows two boxes making a single paragraph and the number of lines reported is two:

one
two
2	

That number is not always reliable as we can see from:

```
\dontleavehmode
```

```
\hbox{\strut one}
\break
\vbox{\strut two\par \strut three}
\par
\the\prevgraf
```

Again we get two, which indicates that a line is a box and not a text line which is why for standard TeX display math a value of three is added, assuming simple single line formulas of not that much height and with spacing around it that make it somewhat similar to three lines.

one
two
three
2

The way these three lines are communicated is via `\prevgraf` which is a property of the currently worked on list. When the par builder kicks in, its current value is taking as starting point for the first line number and line numbers are used to pick up the right indentation (from hang or shape). This means that when the hang or shape parameters are not reset, as in the case of a mid paragraph display formula in TeX, we pick up where we left. So, that made me wonder what happens if we set the `\prevgraf` ourselves? As with `\prevdepth`, which has its own interesting side effects, we can get some surprising outcomes. For showing this we use the following macro:

```
\def\Test#1#2%
{\begingroup
 \hsize 9cm
 \hangafter -3
 \hangindent 1cm
 sample data % from a sample file
 \prevgraf #2
 [\number\prevgraf]
 \par % not needed in ctx
 \endgroup}
```

So let's start with the zero case. We use a quote from Mathematical Writing, a course by Don Knuth that Mikael and I agree we would have failed to pass:

Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by “*blah*” or some other grunting noise.[0]

The output is as we expect: we get three indented lines. So what happens when we set the `\prevgraf` to one?

Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by “*blah*” or some other grunting noise.[1]

We’re good, so let’s try two:

Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by “*blah*” or some other grunting noise.[2]

Going up from there we see that values larger than three seemingly get ignored, so a safe thousand produces:

Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by “*blah*” or some other grunting noise.[1000]

I then became curious what happens with effectively zero values but that is hard because when you set `\prevgraf` negative assignments give an error. That made me wonder what would happen with the maximum positive value, given that I’d seen the code involved doing additions:

Many
read-
ers
will
skim
...
other
grunt-
ing
noise.[2147483647]

Okay, that looks interesting! So what if we go a bit below that?

Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow

```
smoothly
when
...
or
some
other
grunt-
ing
noise.[2147483645]
```

We now come to the question: “*Is this really what Don Knuth had in mind?*” For sure the answer is “*yes*” if we can come up with a use case. However, because once a user knows that this feature is there, and because we cannot be sure (without examining the code) if a wrap around can make \TeX fail, it is best to just catch this case. In order to remain compatible we do this via an option: we set the third bit in `\normalizeparmode`. In ConTeXt speak:

```
\bitwiseiflip \normalizeparmode \limitprevgrafnormalizeparcode
```

As you can see below, we’re fine now.

Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow smoothly when ... or some other grunting noise.[2147483647]

Can we come up with a use case? Sure. In plain \TeX there is a macro that will show the word snippets on the console:

```
\def\showhyphens#1%
{\setbox0\vbox
 {\parfillskip Opt
  \hsize \maxdimen
  \tenrm
  \pretolerance-1
  \tolerance -1
  \hbadness 0
  \showboxdepth 0
  \#1}}
```

Now watch the extreme `\hsize` here. We clip the one line result to the text width. One would expect a reasonable sample text to fit on that line and indeed it does but the tolerance values combined with the lack of a final filler will lead to underfull box messages.

Many

If we use `.2\maxdimen` instead we see this:

Many	readers	will	skim
------	---------	------	------

Here comes our use case. We limit the text because here we get multiple lines. Interestingly, we don't even need to set `\hangindent` to some value.

```
\def\showhyphens#1%
{\setbox0\vbox
 {#1
  \prevgraf 2147483647
 \%showboxdepth 0
 }}
```

or in ConTeXt speak:

```
\def\showhyphens#1{\setbox\scratchbox\vbox{#1\prevgraf\maxcount}}
```

So one can argue that letting `\prevgraf` wrap around is a feature that permits efficient trickery.

Many
read-
ers
will
skim
...
other
grunt-
ing
noise.

It is this particular neat use case that turn this '*wrap around*' side effect into a proper feature, although in ConTeXt LMTX we will likely enable the catch. But the fact that on the one hand setting `\prevgraf` has a check for a negative value, while on the other hand the par builder permits '*wrapping around*' might be an argument for some "*grunting noise*".

post scriptum

In TeX there is a difference between these two cases:

```
some text \par
$$ formula $$ \par
some text \par
```

and

```
some text % stay in par
$$ formula $$ % stay in par
```

some text \par

In the first case we have three paragraphs and if we use some hanging indentation or par shape it only covers the current paragraph. In the second case we stay in the paragraph and therefore the hang or shape applies to all three. Now here comes the interesting bit.

Compare

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

$x = 1$

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

With

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

$x = 1$

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

and

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

$x = 1$

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

As I mentioned before, a display formula counts as three lines and that is communicated to the par builder by adapting \prevgraf. This is actually independent of the amount of spacing we add so you can take it into account when setting up such a hang. It is one of these cases where more clever automatism makes little sense.

Because spacing around formulas can be tricky in ConTeXt we always managed it differently. For that reason in LuaTeX we added some more engine control over the scenarios, for instance in relation to `\baselineskip` related catches. We just default to no spacing at all. In LMTX we emulate display mode even more different and also use the LuaMetaTeX option to make the dollar driven display mode equal to inline, in order to force users to use the proper mechanisms. In ConTeXt we also have some plugins into `\everypar` that (re)set the shapes so in order to produce the above I had to bypass that. I can't speak for other macro packages but in plain a test looks like this:

```
\abovedisplayskip Opt
\belowdisplayskip Opt
\hangindent      1cm
\hangafter       -7
\input {knuthmath} % 3 lines (@12pt)
$$ x = 1 $$      % 3 lines
\input {knuthmath} % 3 lines
```

In this document I cheated by adding:

```
\everypar          {}
\mathdisplaymode  0
\mathdisplayskipmode 0
```

So now we know one reason why `\prevgraf` is really there and why wraparound is not important. Of course one can wonder why that (current list) property can be set at all by a user. Here's one:

Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.
$x = 1$
Many readers will skim over formulas on their first reading of your exposition.
Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by "blah" or some other grunting noise.

How did we get this? The par builder kicks in when we handle display math but it retains the shape. That means that after the formula is done we have a positive value of `\prevgraf` and that means we can subtract from it without going negative (and triggering an error):

```
\input {knuthmath}
$$ x = 1 $$
\prevgraf\numexpr\prevgraf-2\relax % etex
\input {knuthmath}
```

As it is not a (engine or user) register we can't advance and because it operates on the current list grouping is irrelevant, so for pure TeX one has to do something:

```
{\count2=\prevgraf \advance\count2 by -2 \prevgraf=\count2} % tex
```

So, if really needed one can typeset the formula and from its dimensions and spacing above and below derive the correction. Here we just subtract two. There is no way that Mikael can convince me to add a fragile feature like this to ConTeXt, and I know he won't. We anyway have different nuts (or nut math) to crack. Also, we promote a more spacy source setup so users will definitely add empty lines (`\par` equivalents) before and after formulas, which makes all the above fail. But out of curiosity, maybe we should get the latest TAOCP collection and check where display math is used in a hang or shape context.

Below I will collect some of the questions and remarks-turned-questions that keep popping up and start annoying me, especially when they come from people who should know better (being involved in development themselves). I'm always puzzled why these things come up, especially by people who are no user and should not waste time on commenting on ConTeXt.

All these versions, ConTeXt keep changing, so what's next?

Sure, we're now at the third version, MkII, MkIV and LMTX, but there is some progression in this. The first version evolved from \TeX to e- \TeX to pdf \TeX (but also could handle $\text{X}\text{\TeX}$ and Aleph). But in order to get things done better we moved on to Lua \TeX and because that is a ConTeXt related project it made sense to split the code base which made us end up with a frozen stable MkII and an evolving-with-Lua \TeX MkIV. Then there was a demand for a stable Lua \TeX for usage otherwise which in turn lead to the LuaMeta \TeX project and its related ConTeXt evolution LMTX. So, yes, this macro package keeps changing. And it this bad? Don't other macro packages evolve? And why do users of other packages bother anyway? I never heard a ConTeXt user complain either. By the way, how do other macro packages actually count and distinguish versions?

Why is ConTeXt so slow?

Because I seldom hear complaints from users about performance, why do users of other macro packages find reason to even bother. In MkII we immediately started with a high level keyword driven interface so that came with a price. But quite some effort was put into making it as fast and efficient as possible. Fortunately for ConTeXt users the MkIV version became faster over time, in spite of it using a 32 bit engine (which comes at a price). Even better is that LMTX with LuaMeta \TeX has gained a lot over MkIV. But then, I guess, other macro packages that use Lua \TeX are also fast, so maybe the claims that ConTeXt is much slower than other macro packages still hold. I'm not going to check it, and I bet ConTeXt users don't care.

Why does ConTeXt (even) needs a runner.

Indeed, because we don't want users to be bothered with managing runs right from the start it came with a program (Modula2) and later a script (Perl followed up by Ruby) that checks if an additional run is needed because of some change in the table of contents, references, the index, abbreviations, positioning, etc. Index sorting was done too so there was no further dependency. We thought that was actually a good thing. With Lua \TeX and LuaMeta \TeX all that became even more integrated because Lua was used. The runner(s) also made it possible to ship additional scripts without the need for potentially clashing applications in the ever growing \TeX ecosystem. Interesting is that ridiculing ConTeXt for script dependency was never complemented by ridiculing

other macro packages that nowadays seem to depend on scripts (with some even using LuaTeX which originates in the ConTeXt domain).

Why does ConTeXt organizes files that way?

ConTeXt sticks quite well to the TeX Directory Structure, so what is the problem here?. Yes, we needed some granularity (e.g. for MetaPost) but later that just became normal. And indeed we optionally let users use a flat directory structure for fonts but that's normally in the users own local tree. Oh, and in MkIV and LMTX we use our own file database (actually also in MkII at some point), just because (definitely at that time) it was way faster and we needed more features. The same is true for the font database, utf encoded hyphenation patterns, and so on. Can it be that we're often just ahead of the pack?

Let's nor forget to complain about the fact that MkIV and LMTX use a cache but so do lots or programs: just think browsers of some scripting language ecosystems. And that was introduced right after we started with MkIV and hasn't changed much at all. Users expect no less. And other macro packages are free not to use the cache (for e.g. fonts).

The authors of ConTeXt don't care about compatibility, do they?

You're joking, right? Surely some features became sort of obsolete when we moved to MkIV, like encodings. But if users like to stick to them, they can. Do you really think that user like us to drop compatibility? Maybe it fits some narrative to spread that story. Of course, we make things better if we can, and the interfaces have always permitted upgrades and extensions. There are definitely cases when (maybe due to user demand) something new gets added that then evolves towards a stable state, so yes, there can be code in flux. But that is natural. Should we just assume that other macro packages don't evolve, never have bugs, don't break anything, never fix broken things immediately? Maybe. And complaining about ConTeXt evolving is none of its non-users business anyway.

Is ConTeXt commercial?

This is one of the strangest questions (or remarks). We use ConTeXt ourself and using it in a job is by definition commercial use. Are all other TeXies only using TeX macro packages in the free time, as hobby? I'm pretty sure that more money is made by competing package users and I'm also sure that most of the time involved in creating ConTeXt (and LuaMetaTeX for that matter) is not covered by income. Using the fact that ConTeXt is developed by a (small) company excuse for lack of development elsewhere is about as lame as it can get. Much development is done without us needed it, but because we like doing it, because of the challenge.

Should I use ConTeXt for math?

Of course, because that's what TeX is good at. If you are forced to use a specific macro package for its math abilities, just do so. If you want to move on or want consistent

interfaces, maybe ConTeXt is for you. We don't care. Trust your eyes more than assumed standards or ways of doing math typesetting.

Why is the format file so much larger than for other packages?

The answer is simple: we have an integrated system, so we have plenty macros and with each token taking 8 bytes (data and link) that adds up. And for MkIV and LMTX there also Lua code involved as well as a rather large character database. In LuaTeX the format file is compressed (and also zipped) and in LuaMetaTeX is it is a bit more compressed but now zipped; still the LMTX format file is smaller than the MkIV one. We let those who complain wonder why that is. We also let users of other macro packages wonder if loading a ton of stuff later on doesn't accumulate to a similar or larger memory footprint. And, as with many critics: make sure to check every few years if that other macro package hasn't caught up and can be criticized the same way.

Edith Sundqvist
Mikael Sundqvist
Hans Hagen
Willi Egger

Male dominance

When you start using \TeX you can't get around the fact that it uses dimensions. You have to set up a paper size, configure a line width, tell it what font size to use, etc. As with many techniques that evolved in different countries the way to express a dimension can be done differently. In Europe we like to talk in centimeters (**cm**) or millimeters (**mm**) and in the United States it's all about inches (**in**). Typographers all over the world speak in terms of points (**pt**), didots (**dd**), ciceros (**cc**) and picas (**pc**) while those messing around with digital typography prefer "big" (PostScript) points (**bp**). \TeX ies sometimes like scaled points (**sp**) as 1 **sp** is the smallest internal representation of a unit. When someone talks "*points*" you can't be sure if it is big points or \TeX points because the **pt** unit is often used for both.

There are also font-related units, like the popular em width (**em**) and ex height (**ex**) and there is even a pixel unit (**px**) that can be set to some resolution but that one is rarely used. There is also a math unit (**mu**) that scales with the math font in use.

All units are internally scaled points and one real point is 65536 scaled points. That means that when a unit is entered it gets mapped onto this internal scaled point quantity.³⁵

For a while we had the new didot and new cicero but in $\text{LuaMeta}\text{\TeX}$ these were dropped because no one used them. On the contrary, the newly introduced (Don) Knuth unit (**dk**) is quite convenient and we use it as a convenient offset for a so-called \TeX page environment, which we use a lot in testing math functionality.

From this summary we can observe that there are three units that are names: Didot, Cicero and Knuth. But do you realize that these are all males? That can't be right and should be fixed. If you look at user styles (or questions on support platforms) you will also notice that in spite of standardization, the inch (**in**) has not been replaced by its more correct metric counterparts. Okay, that might be due to the fact that there is no meter as unit but using smaller dimensions (**mm** and **cm**) makes more sense, also for

³⁵ If you go back to the early days, there are even cases where you want to talk in terms of true units. Those are not affected by original \TeX 's magnification factor (**mag**) but in $\text{LuaMeta}\text{\TeX}$ we dropped that factor and therefore also these true units became obsolete.

internal accuracy reasons. That said, it is about time that we eradicate the inch or at least come up with something more metric.

So there you have it: we need some female units that correctly stay within the metric domain! In order to convince users to drop the inch the first new unit somewhat relates to it: one Edith (es) is the median of the widths of thumbs of BachoTeX 2023 attendees. One can argue that this is somewhat arbitrary and indeed it is. In order to get a decent value we use a discrete measurement device that groups thumbs into 15, 20, 25 and 30 mm intervals. A 10mm interval is unlikely to get many hits unless the TeX ecosystem suddenly became very easy to use and toddlers get interested in it as a game.

Rule of thumb

If we talk in terms of one Edith, we should keep in mind that at any point we can decide to re-calibrate that unit. If we end up below 25mm we probably have quite some young and/or old users in the sample set. So, in order to have a constant value, the community has to make sure that TeX (and preferably ConTeXt) usage is nicely distributed. Now, of course at BachoTeX we are quite tolerant, because also Plain and L^AT_EX users are sampled. Also, given that this sample of the TeX community is skewed to older users, one can wonder how that influences the initial value. It is up to the ConTeXt group to decide when and where to re-calibrate at a later moment. After all, we have to keep the narrative that ConTeXt is unstable and evolves alive, and occasionally updating a unit fits into that narrative. If you think that this kind of research is somewhat flaky, keep in mind that probably all research related to typography is kind of subjective and somewhat unreal. And BachoTeX being tagged as '*conference*' adds a lot of credibility.

The Edith (—) makes a nice unit for margins, but it is a bit large for offsets, so we also need a female counterpart for the Knuth (■). This is why, just like a centimeter (—) has a smaller companion in millimeter (■), the Edith has a companion Tove (■). In terms of points one Tove is 7.11317pt, while a Knuth is 6.43985pt. It is surely just a coincidence that the value of one Tove in points is about the age of Tove when she became aware that her dad was a ConTeXt fan. In terms of points one Edith is 71.13177pt which, ignoring the unit, comes close to the average age of those who have attended BachoTeX more than 10 times.

The implementation of these units in LuaMetaTeX is not that hard, simply because scanning for these dimensions happens in few places: when scanning dimensions, and in a Lua helper that converts a string to scaled points. At the ConTeXt meeting where we implemented the Knuth, there was some trial and error involved in order to get the right numerator and denominator. One dk is 422042 scaled points which brings us to a numerator 49838 and denominator 7739. Except for scaled points, the fraction gets multiplied by 65536 and the amount. Most units have numerators and denominators with weird values, although 7227 jumps out.

unit	visualized	name	num	denom
bp	—	base point	7227	7200

cc	■	cicero	14856	1157
cm	■■■	centimeter	7227	254
dd	■	didot	1238	1157
dk	■	knuth	49838	7739
es	■■■■■	edith	9176	129
in	■■■■■■■	inch	7227	100
mm	■	millimeter	7227	2540
pc	■■	pica	12	1
pt	■	point	1	1
sp	■	scaled point ³⁶	1	1
ts	■■	tove	4588	645

When you consider these numbers it is good to realize that internally the engine uses a 32 bit number, split into two halves. There is a maximum, 16383.99998pt, so that (intermediate) calculations don't overflow. The last digit of what \TeX reports when it computes a dimension as points is to be taken with a grain of salt. Here is how the Edith and Tove compare to their metric counterparts:

```
2.5cm 4661699 71.13188pt
2.5mm 466169 7.11317pt
1es 4661692 71.13177pt
1ts 466169 7.11317pt
```

In case you wonder if checking for yet another unit has drawbacks in terms of performance, we can guarantee LMT χ users that they won't notice a performance hit. Even with these additional units the engine quite likely beats its predecessors in scanning units. And the impact on the code base is less than 20 short lines of trivial code so that goes unnoticed as well.

Calibration

In order to conduct the calibration we need a reliable measurement device and here we got lucky. The ConTeXt community has some unique craftsmanship amongst its members and Willi Egger made us a robust sampling device that can compete with those used by the ones that the International Organization for Standards uses: the Edithorial.

In addition to that, the ConTeXt Math Society, indeed the same one that brings you all these nice new math capabilities in LuaMetaTeX, provided the necessary statistical and mathematical underpinning to make the Edith and Tove believable units. So here are some more details.

³⁶ This one is not multiplied by 65536.



Figure 21.1 Results from thumb measurements at BachoTeX, with the median thumb marked in blue.

We have found out that the Tove unit, 2.5 millimeters, corresponds to 7.1131744384765625 points. Let us find a decent rational approximation of this, with a small denominator. We do this by calculating the continued fraction, and we try a few steps to get something that is good enough.

We start by noting that the integer part is 7. We then use a calculator (in our case Wolfram Alpha) to compute

$$\frac{1}{7.1131744384765625 - 7} = \frac{1}{0.1131744384765625} \approx 8.835917486854523392207091816098$$

This means that we get as a first possible choice

$$7 + \frac{1}{8} = \frac{57}{8} = 7.125.$$

We continue, and next note that

$$\frac{1}{8.835917486854523392207091816098 - 8} \approx 1.196290322580645161290322580645.$$

Thus, our next candidate is

$$7 + \frac{1}{8 + \frac{1}{1}} = \frac{64}{9} = 7.\bar{1}.$$

Here, the bar over the 1 indicates that 1 is repeating. In the next step we calculate

$$\frac{1}{1.196290322580645161290322580645 - 1} \approx 5.094494658997534921939194741167.$$

The next candidate becomes

$$7 + \frac{1}{8 + \frac{1}{1 + \frac{1}{5}}} = \frac{377}{53} = 7.\overline{1132075471698}.$$

We continue, to get

$$\frac{1}{5.094494658997534921939194741167 - 5} \approx 10.582608695652173913043478260870$$

The next approximant becomes

$$7 + \frac{1}{8 + \frac{1}{1 + \frac{1}{5 + \frac{1}{10}}}} = \frac{3834}{539} = 7.\overline{113172541743970315398886827458256029684601}.$$

For the next step we have

$$\frac{1}{10.582608695652173913043478260870 - 10} \approx 1.716417910447761194029850746269$$

so the next approximant becomes

$$7 + \frac{1}{8 + \frac{1}{1 + \frac{1}{5 + \frac{1}{10 + \frac{1}{1}}}}} = \frac{4211}{592} = 7.1131\overline{756}.$$

Since this one has such a nice short repeating set of decimals, we fell for it, and quit here. The next approximant would be

$$\frac{12256}{1723}, \frac{20301}{2854}, \frac{28346}{3985}, \frac{48647}{6839}, \frac{466169}{65536},$$

where the last one exactly equals what we started with, 7.1131744384765625. Before we continue, we mention that $[7; 8, 1, 5, 10, 1, 1, 2, 1, 1, 9]$ is a more compact way to write the continued fraction above.

One could perhaps first think that multiplying the rational number by 10 would yield a very similar continued fraction, but that is not the case. In fact, the continued fraction for 71.131744384765625 is given by $[71; 7, 1, 1, 2, 3, 1, 3, 1, 3, 1, 3, 2]$. This put us in a bit of an awkward situation. Do we want a nice approximation for the true value, or do we prefer to have **es** to be exactly 10 times as large as **ts**? If we go for the latter, we could take $42110/592$. We calculated the approximants though, and got

$$\begin{aligned} & \frac{489}{7}, \frac{569}{8}, \frac{1067}{15}, \frac{2703}{38}, \frac{9176}{129}, \frac{11879}{167}, \frac{44813}{630}, \\ & \frac{56692}{692}, \frac{214889}{3021}, \frac{271581}{3818}, \frac{1029532}{14475}, \frac{2330845}{32768}. \end{aligned}$$

When we saw this, it was irresistible to define **es** as

$$\frac{9176}{129} = 71.\overline{131782945736434108527}$$

and then to define **ts** as

$$\frac{9176}{1290} = \frac{4588}{645} = 7.1\overline{131782945736434108527}.$$

The edithorial device

The design of the edithorial also involved some research. Of course there was some discussion about the right way to sample thumbs and those who have attended Ba-choTeX and ConTeXt meetings will not be surprised that Willi is responsible for this. He presented us with a drawing (figure 21.2) that we immediately agreed upon.

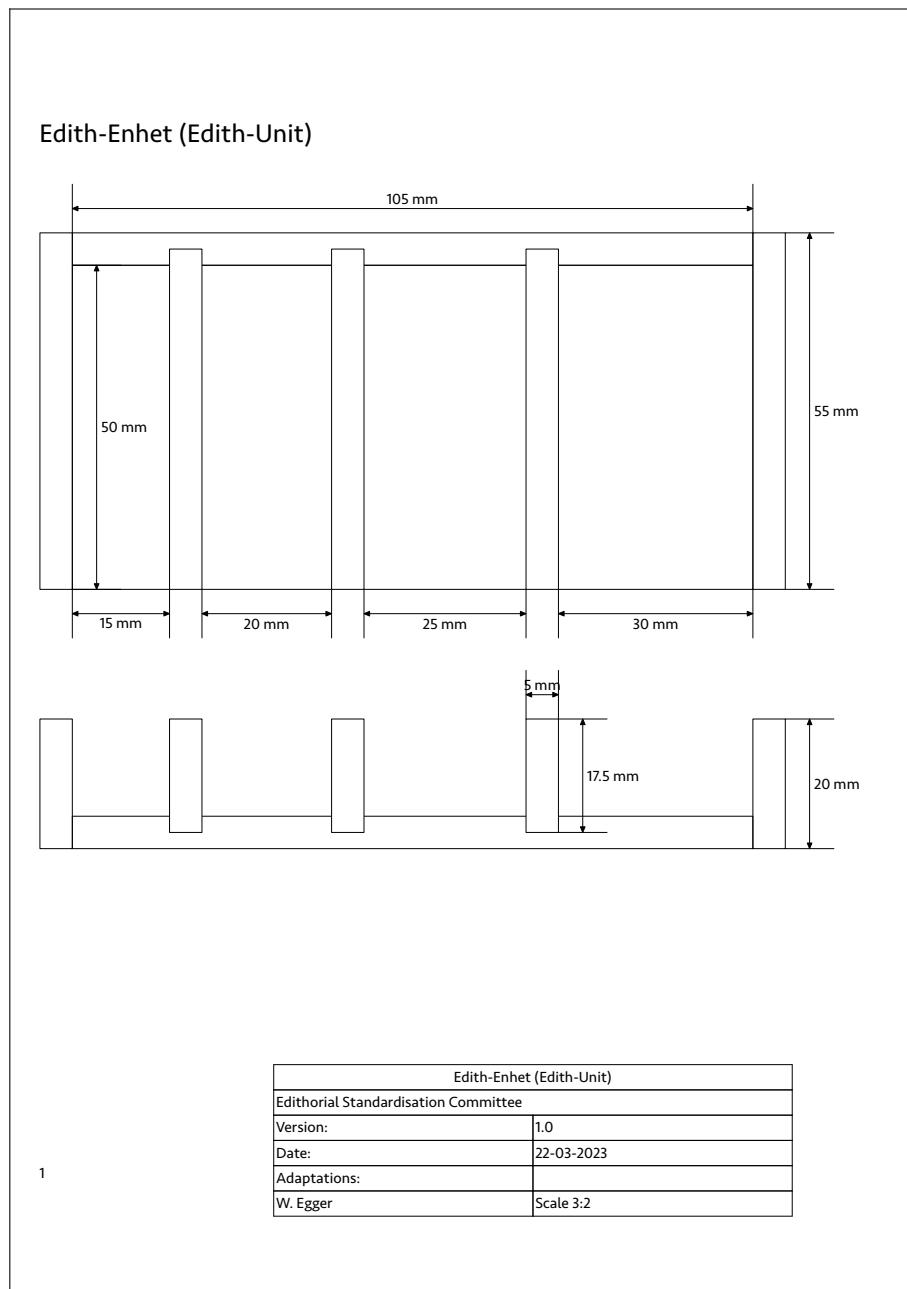


Figure 21.2 A \TeX community worthy edithorial for measuring the Edith.

Willi then sat down and made a prototype (figure ??) in order to see if sampling would work out. Knowing that the device would be stored under harsh conditions in the university city of Lund in Sweden, it had to be sturdy Polish oak and after being brought to precision it underwent first an iron acetate treatment and after that a furniture oil (tung oil) treatment as can be seen in figure 21.3. Even with \TeX being digital we cannot get

around physical devices for measuring digits. And with TeX operating in nanometers we have to fit in.

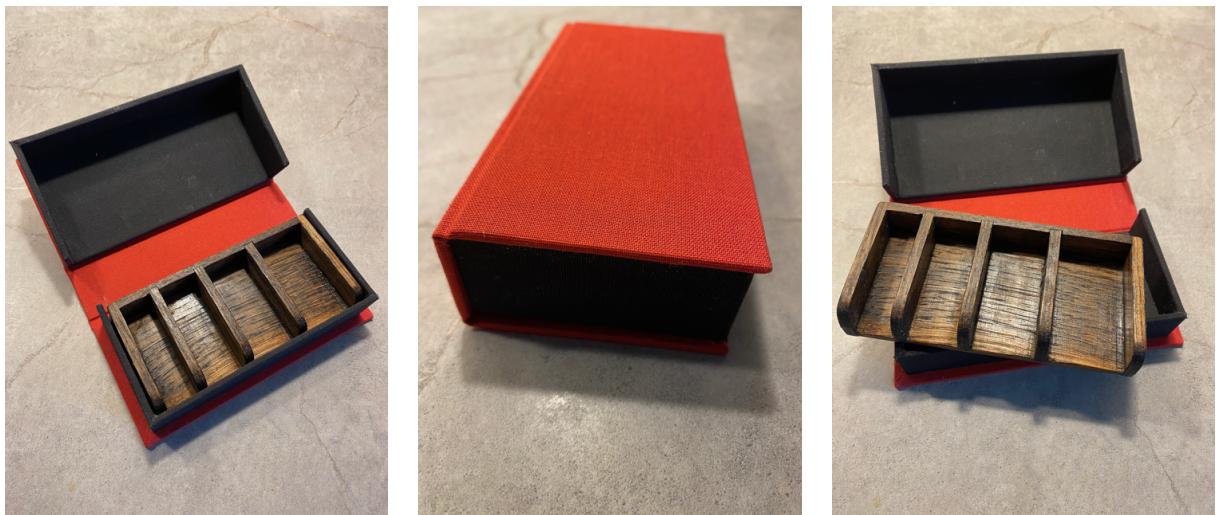


Figure 21.3 This reference editorial plus protective cover.

Some double checking

There is one question we have to answer before we dare to use the Edith (**es**) and Tove (**ts**) as offsets next to a Knuth (**dk**) and that is: in what box does Don's thumb fit? After all, we need to assign some more weight to his thumb. On the Internet you can find images of Don Knuth sitting behind an organ but for reasons of copyright we cannot show these, but one thing we can be sure of is that his thumb is not wider than a key of that instrument, because according to the Wikipedia page [Musical_keyboard](#).³⁷

“Over the last three hundred years, the octave span distance found on historical keyboard instruments (organs, virginals, clavichords, harpsichords, and pianos) has ranged from as little as 125 mm (4.9 in) to as much as 170 mm (6.7 in). Modern piano keyboards ordinarily have an octave span of 164–165 mm (6.5–6.5 in), resulting in the width of black keys averaging 13.7 mm (0.54 in) and white keys about 23.5 mm (0.93 in) at the base, disregarding space between keys.”

This definitely keeps Don's thumb out of the 30mm bucket. When we zoom into these images it seems also unlikely that the thumb will go to the 20mm bucket, but in the end the only one who can answer this is Don Knuth himself. And because he's behind an email firewall we don't dare to ask him. So more research was needed and after a brainstorm session we decided to rely on a public visual that any TeX user should be familiar with: the yearly Christmas lectures.

³⁷ Notice how metric measures win over inches here!

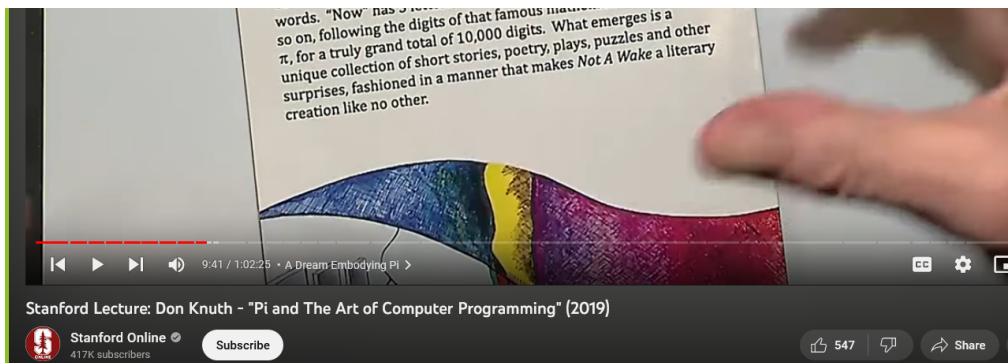


Figure 21.4 The 2019 lecture: Pi and The Art of Computer Programming.

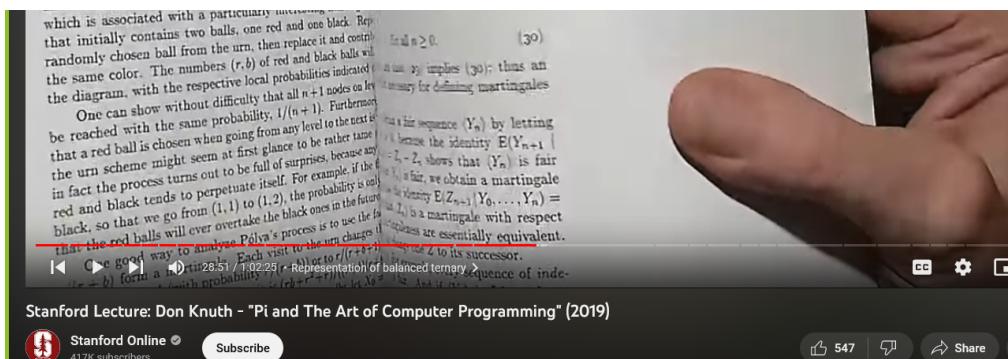


Figure 21.5 The 2014 lecture: (3/2)-ary Trees

And because we know which books the thumb is on, we can calculate the bucket by comparing the dimensions: on one case we use the paper size as a reference, on the other case we use the interline spacing of the book as reference!

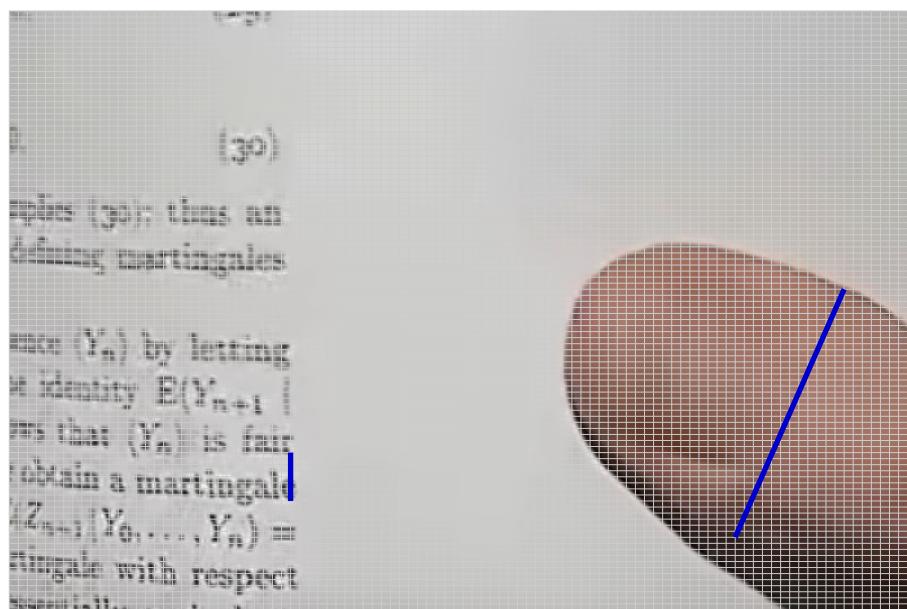


Figure 21.6 Close-up of the thumb.

In Figure 21.6 we show a close-up of the thumb and the page. We have divided the image into a 100×100 grid, but the aspect ratio of the image is $3 : 2$, so we need to

compensate for that. We estimate that the interline space of the text is 8 grid lines high, while the diagonal line measuring the width of the thumb is 12 grid lines wide and 42 grid lines high. This means that the thumb-interline space quotient is given by

$$\frac{\sqrt{(12 \cdot 3/2)^2 + 42^2}}{8} \approx 5.71.$$

Next, we need to know what interline space is used. We should probably know this by heart, but as we do not, we instead downloaded one of the pre-fascicles of TAOCP volume 4. We cut out a square with sides of 5 cm, and added again a 100×100 grid.

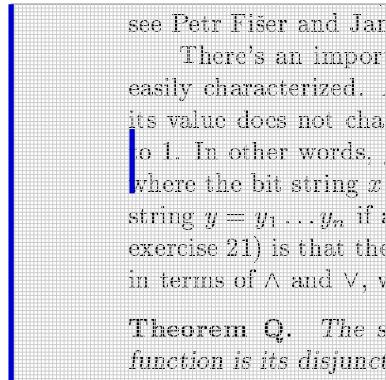


Figure 21.7 A
5cm \times 5cm cutout.

We measured the height of two lines and got in return 17 grid lines. This means that the interline space is given by

$$5 \times \frac{17}{2 \cdot 100} \text{ cm} = 0.425 \text{ cm}.$$

As a result we estimate that Don Knuth's thumb has the size

$$5.71 \times 0.425 \text{ cm} \approx 2.43 \text{ cm}.$$

If we're right about all this then the Edith will not be influenced by the grand wizard's thumb, so the well-calibrated (derived) Tove cannot be discarded for offsets as being less accurate (and stable) as the Knuth.

A modern relative unit

Since \TeX showed up a lot has changed when it comes to computers: the computers considered powerful in the early days now fit in your pocket. One disadvantage of these portable devices is that they have a variety of display sizes. A document can easily be generated again, adapting the layout to all these devices is a bit of a pain.

This is why we introduce a new dynamic unit, the **eu** or the European Unit, but one that can be changed by setting an internal register, **\eufactor**. Because that defaults to 10,

one `eu` starts out as one `es`. A nice coincidence is that one can also read it as Edith's Unit.



We can set the factor in Tove steps between 1 and 50 so that we retain a reasonable accuracy. So, this relative unit stresses the sisterhood of these two new units because `1eu` is `10ts` and `1es`. This unit might also come in handy when writing manuals so you can bet that we will use it.

These units are modern in another way too. The popular game MineCraft has its own unit, a block, as (for instance) discussed on https://minecraft.fandom.com/wiki/Tutorials/Units_of_measure. For those using inches, one inch is 0.0254 blocks, so one block makes 39.3700787in. For those using metric system one cm equals 0.01 blocks or 0.16 pixels and therefore one block makes 40cm. These 40cm are 16 Ediths which means that the Edith is also a good introduction in the hexadecimal numbering system. Unfortunately LEGO bricks are defined in inches so there the inchers still have the edge. But Edith and Tove have an advantage in MineCraft, which is confirmed by observation. Just like some of \TeX 's units are actually defined using the inch paradigm, we could add units like `mb` for MineCraft Block being 16 Ediths. After all, implementing extra units is trivial in LuaMeta \TeX . Let us know what you think.

How about MetaPost?

We not only have to deal with \TeX but also with MetaPost, so from now on MetaFun will also provide these units, which we can then use to properly draw thumbs as in figure 21.8.

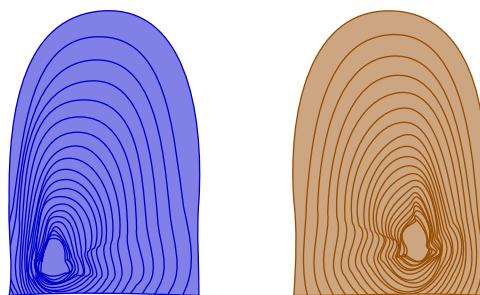


Figure 21.8 One can sign documents with these calibrated thumbs.

While checking other units in MetaFun we were reminded that they are there given as floats and not as fractions. We were amused to see

```
mm := 2.83464 ;
cm := 28.34645 ;
```

which means that a `mm` is not exactly one tenth of a `cm`, and also that the rounding has been done by the even/odd rounding off rule. We decided to define

```
es := 71.13174 ;  
ts := 7.11317 ;
```

Overflow

When you enter a dimension in \TeX and it is larger than 16383.99998pt or 1073741823 scaled points, an error message is shown and when you ask for help, that contains the sentence "*I can't work with sizes bigger than about 19 feet*". There is no `ft` unit in \TeX , so the user has to do some conversion, maybe taking ones own foot into account.

Just like we had to adapt the error message issued when an unknown unit is used, we decided make the overflow message a bit more detailed. For that we introduced the Theodore, where that unit is to the Edith what the Foot is to the Inch. With one Theodore being five Edits we now report this:

"I can't work with sizes bigger than about 19 feet (45 Theodores as of 2023), 575 centimeters, 2300 Toves, 230 Ediths or 16383 points."

So how did we come to this one? At the Bacho \TeX meeting the 18 month old, always good humored, Theodore was running around in the conference room and his little feet were carefully measured by his father Arthur Rosendahl (the self appointed High Commissioner of Hyphenation and upcoming tug president). Because the 19 inches are also an approximation, we rounded the Theodore to five Ediths. In addition we mention a few more maxima, so that the user gets a better impression how large \TeX can go. Mojca Miklavec, who gets her feet dirty by managing the binary build farm on the context garden, proposed a `th` unit but as there is no `ft` we didn't come to a conclusion yet. Actually that unit would make a good default for text width, just like an `es` makes perfect left margin, and a `ts` a nice offset around framed content.

Wrapping up

In this article we discussed a few additional units that have been added to LuaMetaTeX . We've carefully chosen some names that not only compensate the male dominance in unit names, but also have a modern and fresh ring. The units are of course metric. The Edith (`es`) replaces the deprecated inch (`in`) and the Tove (`ts`) can be used for offsets as alternative to the Knuth (`dk`) that of course we will keep using alongside. The units are calibrated using an editorial of which there exists a unique reference measurement piece. The standard has been established at the 2023 Bacho \TeX meeting and might be recalibrated at a future Con \TeX t meeting when a new generation of users thinks that is needed. Many thanks to Karl Berry for copy-editing.

Introduction

When a TeX user posts a math related font question on SE there is a good chance that it catches our eyes. In this case the question³⁸ was about the lack of kerning in OpenType math usage. The person (signature '*dmaxwell*') showed these examples

df *W*, *df* *W*,
pdftex luatex

and asked if it is possible to use kern pairs also in OpenTypemath.

Before we dive into the details we show how we can solve it in ConTeXt lmtx. We do in fact have support for kern pairs, and one way to set them up is via the goodie file system. The goodie file for Latin Modern Math is `modern-math.lfg`. There are plenty of tweaks, and we can add an instance of the “`kernpairs`” tweak.

```
{\n    tweak = "kernpairs",\n    list = {\n        [0x1D451] = { -- italic d\n            [0x1D453] = -.2, -- italic f\n        },\n    },\n}.
```

This moves the italic f to the left (the minus sign), and the amount is 20% of the width of the italic d. With a similar entry for the combination of italic W and the comma, we get

df W ,

The kerning here is not as aggressive as the (rather famous) *df* kerning in Computer modern. When we came to discuss this issue, we realized that there were a bit more for us to do, and more to say on the subject that also could interest others, so keep on reading.

³⁸ <https://tex.stackexchange.com/q/692788/52406>

Kerns in OpenType math

Scanning the OpenType math font specification³⁹ for kerning, one might become surprised. Somewhere at the beginning one can read the following

Layout of math formulas is quite different from regular text layout that is done using tables such as GSUB and GPOS. Regular text layout mainly deals with a line of text, often formatted with a single font. In this situation, actions such as contextual substitution or kerning can be done with access to the complete context of the line of text, and the rules can be expressed in terms of known glyph sequences. Math layout is quite different from this.

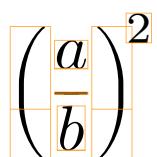
Reading this, one gets the impression that there is no support for kerning in math, at least not pairwise between characters. But further down, one finds that there is indeed support for kerning in math, but the only kerning that is described is what we call staircase kerning, between base characters and sub-, super- and prescripts. They explain it like this,

For any given glyph, different values can be specified for four corner positions — top-right, to-left [*sic*], etc. — allowing for different kerning adjuments according to whether the glyph occurs as a subscript, a superscript, a base being kerned with a subscript, or a base being kerned with a superscript.

and then follows a rather long and technical description of how this is meant to work.

As ‘*dmaxwell*’, asking the question, noticed, this lack of kerning of math fonts might be a mistake. This holds in particular for fonts like Latin Modern Math, that originate from a math fonts that indeed have pairwise kerning.

Even if not specified by Microsoft, \TeX engines can implement additional support. This was done for example in LuaTeX and LuaMetaTeX . In the latter, there is indeed support for the built in staircase kerns, corner kerns, as well as pairwise kerns. In fact, the staircase kerns has proven to not be too useful, so they are automatically transformed into corner kerns (that also work between for example a base character and a superscript, as in the example below).



In the rest of this article we will focus on the pairwise kerns.

Kerns in traditional math

Traditional \TeX font handling involves tfm files. These files provide properties of glyphs. For each glyph it specifies a width, height, depth and italic correction, but

³⁹ <https://learn.microsoft.com/en-us/typography/opentype/spec/math>

does that in a compact way: choosing from a limited set. For the discussion here this is not relevant, apart maybe from the observation that the combination of width and italic correction in math is kind of special. It also explains why there can be differences in metrics between OpenType fonts and its ancestors, just because in OpenType we can be more granular. Optionally a character can have a pointer to a larger size, an extensible recipe, for instance for delimiters. Also optional is ligature and kern information. For efficiency reasons these properties are packed into a table with flags in the entries indicating what we're dealing with.

You can use the `tftopl` program to see what is actually in a tfm file, and here is such an entry for the `N`:

```
(CHARACTER C N
  (CHARWD R 0.8034725)
  (CHARHT R 0.683332)
  (CHARIC R 0.109027)
  (COMMENT
    (KRN C = R -0.083334)
    (KRN C = R -0.027779)
    (KRN C ; R -0.055555)
    (KRN C : R -0.055555)
    (KRN O 177 R 0.083336)
  )
)
```

If we look at the Lua representation that we have in ConTeXt we get this:

```
t={
  ["depth"]=0,
  ["height"]=447828,
  ["italic"]=71451,
  ["kerns"]={
    [58]=-36409,
    [59]=-36409,
    [61]=-54614,
    [127]=54615,
  },
  ["remainder"]=0,
  ["tag"]=1,
  ["width"]=526563,
}
```

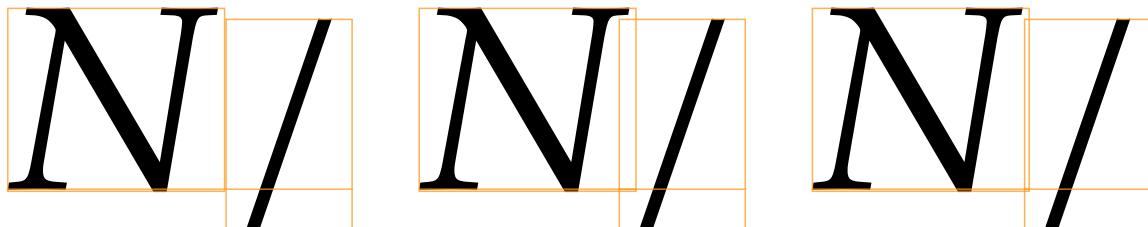
Here we are only interested in the `kerns` sub table and when you compare both listings you will notice that the first one has one entry more and we admit that this caught us by surprise. We were interested in these values because we wanted to add them to the Latin Modern Math goodie file, so what do we have to choose?

When we load this font in ConTeXt, we get a message:

```
duplicate between index 0x4E and 0x3D: -0.83334pt -> -0.27779pt
duplicate between index 0x58 and 0x3D: -0.83334pt -> -0.27779pt
```

The left characters are the **N** and **X**, the right one is the **/**, which is actually one that we are less interested in because we have fraction spacing etc. But that doesn't mean that we wanted to explore it a bit more. The next code shows some variants, we wrap into an `ord` to make sure that no inter-atom spacing kicks in.

```
$\mathord{N\kern 0pt      /}$  
$\mathord{N\kern-0.83334pt/}$  
$\mathord{N\kern-0.27779pt/}$
```



One can argue that the second variant looks best and that is indeed the one that Donald Knuth expects the engine to use. When (in traditional mode) a kern is needed the engine will run over the characters property array and when it hits the kern asked for it quits scanning. That also means that when we use a Lua approach, where we use a hash table instead, we should now overwrite already set entries when the tfm table gets parsed and its data gets stored. Just to be sure we checked all engines, and all behave as expected with respect to applying the first kern.

A natural follow up question is "*where does this duplicate entry come from*". For that we need to go back to the source. Just for the record, the Latin Modern fonts as well as its Computer Modern predecessors behave the same here.

```
ligtable "N":"X": slash kern -1.5u#, % $N:X:$
"C":"T": slash kern -.5u#, comma kern -u#, period kern -u#, %
$C:T:$
"B":"E":"G":"O":"Q":"R":oct"174": % $B:E:G:O:Q:R:\jmath:$
"l":"p":"q":"t":"w":oct"002":oct"004": % $l:p:q:t:w:\it\Theta:\Xi:$
oct"006":oct"010":oct"012":oct"014": % $\it\Sigma:\Phi:\Omega:\beta:$
oct"020":oct"022":oct"032":oct"036": % $\zeta:\theta:\rho:\phi:$
oct"042":oct"043":oct"045": % $\varepsilon:\vartheta:\varrho:$
oct"046":oct"047":oct"100": % $\varsigma:\varphi:\partial:$
skewchar kern 3skew#;
```

This blob of MetaFont code can be found in `mathit.mf` and in order to understand it better we can reformat it a bit:

```
ligtable
```

```

"N": "X":
    slash kern -1.5u#,
"C": "T":
    slash kern -.5u#,
    comma kern -u#,
    period kern -u#,
"B": "E": "G": "O": "Q": "R": oct"174":
    "l": "p": "q": "t": "w": oct"002": oct"004":
    oct"006": oct"010": oct"012": oct"014":
    oct"020": oct"022": oct"032": oct"036":
    oct"042": oct"043": oct"045":
    oct"046": oct"047": oct"100":
    skewchar kern 3skew#;

```

The second `slash` entry will result in the extra entry in the tfm file but as mentioned it will be ignored by TeX. One can argue that `tftopl` should warn for this duplicate. This is the only entry that saves on the skewchar definition, so it could have been like this:

```

ligitable
"N": "X":
    slash kern -1.5u#,
    comma kern -u#,
    period kern -u#,
    skewchar kern 3skew#;
ligitable
"C": "T":
    slash kern -.5u#,
    comma kern -u#,
    period kern -u#,
"B": "E": "G": "O": "Q": "R": oct"174":
    "l": "p": "q": "t": "w": oct"002": oct"004":
    oct"006": oct"010": oct"012": oct"014":
    oct"020": oct"022": oct"032": oct"036":
    oct"042": oct"043": oct"045":
    oct"046": oct"047": oct"100":
    skewchar kern 3skew#;

```

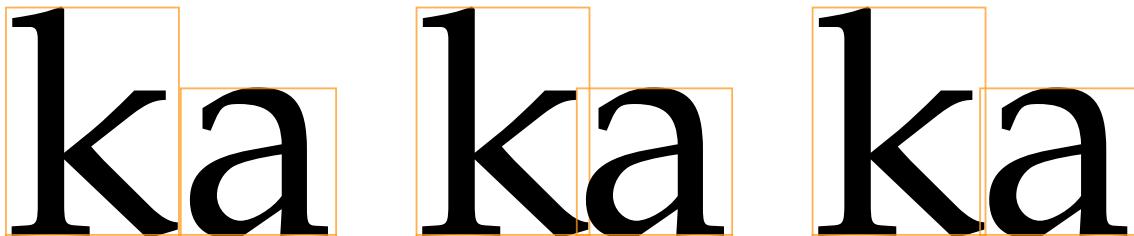
Which then gives just one kern pair for `N` and `X` with `/`, not that anyone ever sees this as a real problem. Of course we were curious if other files in the collection have these duplicate entries and there is indeed one other case, the upright text specification (including slanted and bold but not italic) has this:

```

duplicate between index 0x6B and 0x61: -0.55556pt -> -0.27779pt
duplicate between index 0x76 and 0x61: -0.55556pt -> -0.27779pt

```

Here the second kern is actually better, so one can wonder what the intended kerning was. It might be the reason why the Dutch word 'kaas' comes out so badly in a traditional setup.



It anyhow demonstrates that one should be careful with the order of definitions, and adhere to the specification that the first kern wins. This is actually not that different from the way OpenType works, although, especially in the early days, one could find fonts that were set up with the assumption that the last one wins or that kerns accumulate. Of course much can get unnoticed on a relatively low resolution output medium.

How to proceed?

We now have to decide how to proceed in the perspective of ConTeXt LMTX, where we follow a somewhat different approach to math spacing. As discussed elsewhere, we already rid ourselves of italic correction as a means to anchor subscripts: we went natural width with bottom kerns.

We also have, as shown at the beginning of this article, a dedicated mechanism for math kerns that compensates for the lack of such feature in OpenType math.

Because an OpenType math font is a hybrid font that has upright, italic and bold shapes in one font, kerning can be more extensive than in a traditional assembly. Now, if only the fonts had a kern table, we could just use that one but the fonts lack such a table. For Latin Modern we can add definitions taken from the eight bit fonts but for instance Pagella and Bonum need their own. We use this example:

```
$df + dj + dl+ ab + if + ij$
```

In untweaked Latin Modern we get this

and when tweak we get:

The TeXGyre Pagella and Bonum math fonts give:

$df + dj + dl + ab + if + ij$

$df + dj + dl + ab + if + ij$

However, in the tweaked Bonum we get this:

$df + dj + dl + ab + if + ij$

We consider it an interesting coincidence that we actually already had added math kerns to the goodie file that do something similar to the type df and dj . We also admit that the choice for what to kern is mostly determined by what one encounters and we're sure that this is also the case the limited set of kerns in math italic characters of Latin Modern: there is a reason why dj gets a kern and ij doesn't.

But the main reason why we have these kerns in Bonum is not that we wanted to get better kerning but that because we compensate the excessive '*sticking out at the left*' that we compensate has to be partially decompensated in some cases. But we anyway had to translate the kerns into a fraction of the width (in our case the width plus discarded italic correction) because that is what we use in the goodie files where we need to avoid hard coded dimensions.

This is how Pagella comes out tweaked:

$df + dj + dl + ab + if + ij$

So where can we get the kerns from? Playing a bit with values from the text font alphabets give some cue where we should kern but the amounts can't be used: text and math characters sit differently in their bounding boxes. Using side bearing can help but is also not optimal, So, after half a day playing with this we gave up on that. Instead we now use a feature that starts from the values that we find in Computer Modern set, assuming that italic shapes are similar. These are the most important candidates that get kerning with periods and commas:

$C F H J K M N P S T U V W X Y Z f j r \Gamma \Pi \Psi \delta \nu \sigma \tau$

In the Computer Modern fonts the skew character is used for anchoring accents and in OpenType that is done using the `topanchor` property. So, we can discard these. The same is true for the slash: either we set kerns for all characters that matter, or we don't. For instance Latin Modern has a kern after a slash before A, M, N, Y, Z and Δ , but not before e.g. V and W . Also, if we compare the values for italic correction in Latin Modern Math with Computer Modern, we notice subtle differences in the width-italic ratio: 0.367 versus 0.381 so we cannot be too picky. For now have commented the values but they can easily become a feature.

(Optional) side note

It is no secret that eight bit pdf \TeX beats the Unicode engines when it comes to processing (utf8) input and rendering text. Among the Con \TeX t font loader processing mode, *base* mode comes close to the traditional approach, while *node* mode is slower because all is done in Lua. Node modes is needed when we go beyond simple pairwise kerns and simple (non contextual) ligatures.

In Lua we usually store data in a hashed table and that is also true for kerns in fonts that get loaded. When in node mode, we keep all at the Lua end but in base mode we pass the kerns to \TeX . On previous pages we indicated that properties like kerns get stored as a sequential lookup. Because Lua randomizes the hash using a different seed each run. The order in which kerns end up in \TeX is also random. The way we store information today (like in these Lua hashes) was not an option in the early days of computing which makes \TeX even more of a masterpiece of efficiency.

If we use `texnansi` encoding we have kern pairs between `va` and `vo` and in Latin Modern the first one is the first entry and the second one the sixteenth. So does that have some effect on runtime? A somewhat over the top test using these combinations shows that in pdf \TeX that stored the kerns in original order indeed the `va` kerns takes less time to resolve. The further down the list the kern is, the more time it will cost. This is comparable with for instance access to path points in MetaPost: the assumption is that paths are short when this is needed as looping over a 5000 point path involved quite a bit traversing. This has never been an issue for \TeX because fonts are limited to 256 glyphs which means that of the many accented variants of vowels only a few end up in a specific encoding. Of course the more random order in the Lua hash variant makes for a constant look up time in node mode and a more averaging time in base mode. A side effect is that one can add one own additional kerns in node mode.

Isn't it amazing to what observations a simple question on SE bring up? There is one more thing worth mentioning. The lack of these kerns went unnoticed for quite a while. This can indicate a that either users don't care, or just don't consider it needed, or have adapted by adding manuals kerns. It definitely has not resulted in it being seen as a necessity in OpenType math. It is kind of interesting that a power user was told about this he remarked that checking the period of comma after a capital in math mode was part of some quality assurance process because it could indicate that authors messed the spacing (in MathML and of AsciiMath input). So, adding these kerns already paid back. It also shows that \TeX can adapt very well to whatever circumstances it has to deal with.

23.1 Introduction

Among the first macros we cooked up are the ones that we used to typeset chemistry. They come under the `ppchTeX`, and it was mostly done by Hans, Ton and Tobias. Over time rendering structures by `PiCTeX` got replaced by pstricks and later MetaPost. When we moved to MkIV all that got reimplemented, this time by Hans and Alan and we stuck to MetaPost. We actually also updated the syntax and, because old documents were supposed to be rendered by MkII we took the liberty to sacrifice some compatibility for a more complete and consistent feature set. However we never came to updating the existing manual and the new one became work in progress. We didn't really needed the code and Alan, who did need it, knew what was there anyway.

Then Mikael had some questions and it pushed me into manual mode: upgrade the old one and combine it with the pending new one. That also made me wonder if we could now benefit from some new features in the math engine, which in turn could simplify the code base. This is typically something that takes weeks instead of hours so better do it right from the start. And, while at it, one then of course ends up in looking (again) at arrows (or more specifically: stackers) and from that to Unicode, which in turn is good for some introspection.

23.2 Implementation

Among the (old) complications in dealing with chemistry are the following:

- Simple inline chemical formulas (snippets) of the kind $^{238}_{92}\text{U}^+$ where scripts need to be properly vertically aligned and an upright font is used.
- In running text one wants to use $A + B \longrightarrow C$ but also $A + B \longleftrightarrow C$ or $A + B \iff C$.
- Even more symbolic representations might be on the wish list, think of $A - B = C \equiv D$.

There is more but these examples demonstrate a few features: we need prescripts, proper spacing, line breaks at preferred spots, and some symbols.

- And of course one also wants to annotate like $A + B - > \text{here}C$ but also $A + B < - > \text{there}C$ or $A + B <= > \text{whereevery}C$ where the arrows stretch with the text.
- It would be nice if we can also use the advanced alignment options that are available in math but discussing this is beyond this musing.

23.3 Unicode

Implementing the above has never been that hard but just became a little easier in LMTX. However, when doing that I wondered if there were more than the already present symbols to be taken care of. And so I did a search on the internet for “*unicode and chemistry*”. One of the first hits was “*Five symbols used in chemistry L2/23-193*”, a request for some more arrows. One can search the web for it and see if it is still around. When staring at I wondered a bit about the descriptions:

BALANCED EQUILIBRIUM ARROW

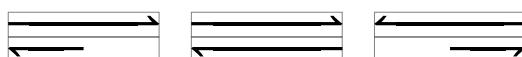
EQUILIBRIUM ARROW LYING TOWARD THE RIGHT

EQUILIBRIUM ARROW LYING TOWARD THE LEFT

REACTION DOES NOT PROCEED

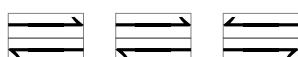
STANDARD STATE SYMBOL

When later I discussed this with Mikael we came to the conclusion that ‘LYING’ probably means ‘LEANING’ but we’re not chemists so we can be wrong. The proposed rendering of the first three boils down to arrows with half a tip, also known as harpoons: long left and right pointing ones stacked for the first and long over short ones for the other two.

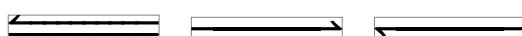


There are several observation to make here. First of all, there is a whole bunch of arrows (stacked and single) that bare descriptions mentioning them being arrows. There is no meaning in them. One can even wonder why some are there. So, in order to be in sync with that it makes more sense to add a few more harpoons. Cooking up names like this serves no purpose.

If we look at the proposed shapes (which are actually different from those used in \TeX packages that are references) one can actually wonder about the way these are supposed to stretch. The short variant is not different from existing double harpoons in which case the meaning is lost. Then when we go longer we get this empty space and then we should wonder how the extensible should kick in: how do the left- and right-most fixed glyphs and the one or two middle repeated ones behave?



Why don’t we just use the next three? After all, as with math, how we interpret symbols depends on how we define them to be read. Sometimes Mikael and I get a good laugh over some of the shapes bound to math code points and we’re sure that some legend is needed when using them.



Of course there are and can be conventions but getting some agreement is not trivial. We only have to look at Unicode math to see some issues, like:

- The repertoire of symbols is large but to a large extent somewhat arbitrary: what was known (or assumed) to be used got in there. There is a difference between a left, right and even middle bar, just like there is a left and right brace.
- There are alphabets but they come with holes. We're not supposed to distinguish between a Plank constant ' h ' and variable ' h '. We can distinguish the greek upper-case \mathbf{A} from a latin \mathbf{A} so mathematicians are less convincing.
- One can define extensible fences with components but not all of them. There are no constructors for arrows. Some likely come from old encodings but why add these and not be consistent.

Plugging a hole in an alphabet is doable. Adding some missing snippets for fences is also no big deal, just because we have a limited set. Making bars consistent also is not hard. But I don't see it happen soon. After all, we're decades along the road and no one bothered much about it till now, and above all, it will never be complete. I wonder if there ever has been a good analysis, extensive description and a watertight arguing from the \TeX community about what should go in Unicode and fonts. It was easy to point to existing fonts and the names used in macro packages and that shows.

To the above we can add:

- There are some combinations of arrows and e.g. rules missing that could have made adding composed constructs easier when there is demand for that. After all, this is what existing characters are also used for. And these chemistry ones fit into this.

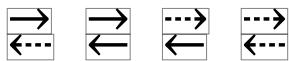
Our solution is to just accept that the \TeX community got what it deserved: a bit of chaos, unreliable cur'n'paste support, symbols but no meaning, imperfect Unicode coverage and therefore imperfect coverage in fonts. The good news is that we can adapt. But we have to be honest: it is not perfect.

23.4 How about

As mentioned we have a problem with base glyphs versus extensibles and these combined hooked things are bad for that. In the chemistry macros the first four are the ones that we always had and the last two are examples of how we can render the two leaning extras.

$\textcolor{brown}{<->}$	$A + B \longleftrightarrow C$	$\textcolor{brown}{<->}$	$A + B \longleftrightarrow C$
$\textcolor{brown}{->}$	$A + B \longrightarrow C$	$\textcolor{brown}{-->}$	$A + B \longrightarrow C$
$\textcolor{brown}{<-}$	$A + B \longleftarrow C$	$\textcolor{brown}{<--}$	$A + B \longleftarrow C$
$\textcolor{brown}{<=>}$	$A + B \iff C$	$\textcolor{brown}{<=>}$	$A + B \iff C$
<hr/>			
$\textcolor{brown}{=>}$	$A + B \longleftarrow C$	$\textcolor{brown}{==>}$	$A + B \longleftarrow C$
$\textcolor{brown}{<=}$	$A + B \longleftarrow C$	$\textcolor{brown}{<==}$	$A + B \longleftarrow C$

Watch the longer should arrow. Not all fonts have it as extensible, which indicates that not that much thought has been put in usage patterns. That said, more interesting would be to use something like this, where we fake a character from two existing arrows:



However, not all math fonts provide these dashed arrows and therefore here we have to use Cambria (here). And even if it has such dashed hashes, they don't come with an extensible. I just want to point out that by looking at what we already have, it might make more sense to extend some (combinations) of those.

It is good to notice that these arrows come with qualifications like “*code for undetermined script*”, so one can wonder how much they relate to math. Actually the fact that we have these holes in alphabets already indicates that math is not really seen as script. Occasionally (old) scripts get added and they get lots of code points, while one can argue for sharing there too, but maybe their status is higher than the status of math. In OpenType fonts math is seen as script but that's because it is basically a selector. One can actually argue for chemistry as a math language in fonts.

The main point I want to make here is that adding some new symbol that is somewhere used but never made it in Unicode in the first place needs some thought. Especially when used in a setting of formulas, where size matters.

23.5 Also

In ConTeXt we have stackers: text above and/or below an extensible, or extensibles above and/or below text. The mentioned arrows are using this mechanism. However, when playing with `\iff`, `\implies` and `\impliesby` in math mode, we noticed some spacing side effects. Originally these double arrows got skips around them but that interfered with our alignment mechanism. However, that could be solved without much hassle by letting the commands check the nature (class) of the previous atom (an indication of being at the start of a next alignment line). In the end that we decided that an extra ‘*implication*’ class was more flexible than adding glue. More interesting was the observation that in Latin Modern and some other math fonts these arrows have different dimensions, which leads to yet another alignment issue.

Again that could be solved by taking the usage into account but one can wonder why the opportunity was lost to make the glyphs consistent with each other, read: come up with a proper analysis of requirements based on decades of TeX usage. At least there could have been recommended alternates. But wait, aren't alternates kind of bad as they demand user intervention (choices)? Sure, but there are more examples of alternates, take the ‘*i*’ (dotless *i*). It has a textual code point but is not in the math alphabets, so one needs alternates as way out (and yes, fonts then have a blackboard and fraktur dotless *i*).

One can argue that these are visual aspects, but with arrows as well as symbols, we have ended up in a somewhat curious inconsistent situation: there are no established command names for the single arrows, so there we speak `\..arrow..` while for some there are names, like `\iff`. The same is true for characters like the dotless *i* and *j*. Some mathematicians use these in the same way as some hard- and software vendors put an

'i' in front of a product name, but in order to get it one has to communicate in terms of `\dotlessi`. or something with an `i` in the macro name.

23.6 Conclusion

To come back to updating chemistry. It makes no sense to add much more. Implementing the left- and right leaning is easy with existing hooks so this is what we will do. Proper math fonts have these and how likely is it that existing math fonts get new ones? If it ever comes to more chemistry in Unicode a handful of them will not help much. It is probably not that hard to find an existing symbol that can act as standard state symbol. In fact, some code points have several additional descriptions so we could just add some to existing one. It's not like we have gone over the top with doing that for math yet.

24.1 Introduction

Having played with and being interested in electronics since I was kid, it will be no surprise that I also keep an eye on RiscV. And because of that I already wondered for a while how easy it would be to compile LuaMeta \TeX on that architecture. Although there are now some decent small boards that can be used, I don't have these, which made me check out if one could emulate it.

24.2 Virtual machines

Back in the days I used VMWare to run a linux machine on my Windows laptop, but nowadays the Linux Subsystem is pretty good for what I want. On the servers we always used (OpenSUSE) linux although at some point we switched from bare metal to virtual machines on top of Proxmox. So, there is some experience with virtual machinery. Of course a question to be asked is "*What does it do with performance?*" but plenty about that has been said in the past so I stick to a summary.

Given enough memory and a decent disk setup, for us running \TeX in a virtual machine is not that significant slower as on bare metal. For instance, running a NAS on the same virtual host is quite efficient, because of the more or less direct connection, while using the same NFS setup over different machine comes at a higher price. The same is true for compiling in the Linux Subsystem: a single core \TeX process can benefit from the other core being used by the host for file management.

Performance of \TeX depends not only on the cpu. Just compare making a format file on a freshly restarted computer with a second make: it's way faster due to the operating system caching the many files that get loaded. The same is true for making the font database although that happens seldom. The console is actually one of the factors that has a high influence.

However, in all these cases we use a real cpu for the task. In the case of this test with RiscV I had to use an emulator (QEMU) so I knew that performance would be kind of bad. Of course that's relative because even (say) 20 times slower would just bring us back to the beginning of this century and it's not like performance was a show stopper then.

24.3 Compiling LuaMeta \TeX

I used my linux laptop that has the same specification as the Windows one: a refurbished Dell Precision 7520 with plenty memory, a fast large ssd, and mobile XEON cpu's.

The linux machine runs OpenSUSE and I was able to install a RiscV virtual machine using QEMU, once I'd figured out how to boot one. It's a bit of a hassle but in the end it works. Compiling LuaMetaTeX boils down to:

- installing the gcc compiler (in this case version 13)
- installing cmake
- fetching the LuaMetaTeX git repository (I used `work`)
- running the `build.sh` script

When we had to support the Raspberry Pi and the Apple M* architectures we had to adapt a few compiler flags and have some specific platform detection (not that dramatic). So I expected that I had to tune for RiscV but in the end the code just compiled. It took however quite some time. Later I gave the virtual machine four cores and then compilation time was sort of acceptable. So, I had a binary!

24.4 Running TeX

Next comes a challenge: how to test LMTX when there is no installer. I mounted my main disk to the virtual machine and just copied the TeX tree from there. Then I copied the binary from `build/native` to the `texmf-linux-64/bin` path and voila, after extending the `PATH` variable I could make a format. Because TeX is a single core process, one has to be patient.

RiscV	17.8 s
linux	1.0 s
Windows	2.0 s

The winner is linux. The Windows binary is a cross compiled MingW one and I think that the difference in performance is mainly due to file handling. In all tree cases the log is piped to a file, otherwise these times are quite a bit higher.

Next I ran the following file, which results in 223 pages:

```
\starttext
\dorecurse {1000} {
    \samplefile{tufte}\par
}
\stoptext
```

This time we need:

RiscV	46.1 s
linux	2.9 s
Windows	2.8 s

The reason why linux and Windows are the same is that there is less file access. We load the format file, some fonts, the main file, and output a few hundred pages pdf.

The emulated RiscV machine is about 15 times slower so not really an option for production but that was not the objective anyway. Actually, if I remember right, the first Raspberry's were not that fast either and those were bare metal and currently most RiscV boards are in its version 3 performance range. Also, this test is mostly done to see what is needed to support RiscV: nothing special.

24.5 Conclusion

The main conclusion is that it was actually quite easy to support this new architecture. Of course I now need to come up with an installer but I delay that till I have hardware because buying some just for testing makes little sense. After all I don't have an M3 for testing either.

The installation script has been extended with a `--native` option (one can also set that in the installer file) in which case files will be updated. This assumes that the user has compiled a binary and put it in the installations `bin` path alongside the `mtxrun.lua` file there. The user also has to make sure that `tex/texmf-native` has the updated binary. It's a bit more work but at least one has an updated system.

In the end a 15 time slowdown is not that dramatic. We've made ConTeXt a bit faster with LuaMetaTeX anyway. My best laptop will be 1.5 times (or more) faster and I expect upcoming RiscV boards to catch up too, just like all these arm machines did.

Just for the record: I didn't try to compile LuaTeX. That source tree is much larger and there are all kind of dependencies involved. It takes way more time to compile anyway. But I expect it to work out well too.

25.1 Introduction

In early May 2024 Karl Berry forwarded me a mail by Nelson Beebe referring to a discussion on a font forum where surprising ligature replacement in OpenType fonts was mentioned in the perspective of abuse: fonts as programs. Anyone familiar with this technology will surely confirm that ligatures are a feature that can be abused, but also that in practice it seldom happens. So let's explore this a bit. Before I show some examples I'll summarize how \TeX deals with ligatures. I'll also make clear that it is unlikely that ligatures are used for unexpected replacements.

This article appeared in a 2024 TugBoat. As usual, thanks to Karl for making it better!

25.2 Ligatures

When \TeX adds characters to the horizontal list it can replace successive character pairs into a single character, a so-called ligature. We end up with a mix of character and ligature nodes. When the paragraph is broken into lines hyphenation is applied when it makes sense which in turn means that these ligatures (and kerns) are temporarily discarded for the word being checked. For that purpose the node that stores a ligature keeps track of the characters that made the ligature. Although it is not that relevant for what we discuss here, it is good to keep in mind that whatever sequence of characters we convert to another sequence of characters, hyphenation can be an influence and whatever we intend might not work well at a line break.

In \LaTeX we don't have character and ligature nodes but only glyph nodes that can be either a normal character or a ligature. We (can) still provide the components that made the ligature but in \LaTeX we don't have these. If needed, the subtype can be used to distinguish between a character and ligature. We don't need the components because hyphenation has already happened before ligatures and kerns are applied.⁴⁰

Although there are these differences between engines in how and when they construct (and in traditional \TeX deconstruct) ligatures, in both cases they are the same: a multiple to singular mapping. The \TeX engine looks at two characters. For example, an 'f' and 'i' become an 'fi' and an 'f' followed by a 'f' becomes 'ff'. Of course that only happens when a font has such a ligature. The new character 'ff' itself can combine with a following 'i' into a 'ffi'. So, even for a three character ligature we go via dual replacement, so for more complex ligatures we need intermediate ligatures or bogus characters.

Deep down \TeX has several kinds of ligatures (the following description is taken from the \LaTeX manual). When \TeX inserts a new ligature, it puts the new glyph in the

⁴⁰ An exception is Arabic where marks are anchored to shape ligatures.

middle of the left and right glyphs. The original left and right glyphs can optionally be retained, and when at least one of them is kept, it is also possible to move the new '*insertion point*' forward one or two places. The glyph that ends up to the right of the insertion point will become the next '*left*'. In this table the **|** indicates the final insertion point.

case	characters	ligature	outcome
0	a b	c	 c
1	a b	c	cb
2	a b	c	ac
3	a b	c	acb
4	a b	c	c b
5	a b	c	a c
6	a b	c	a bc
7	a b	c	ab c

In the first four cases the current position (pointer in the node list) moves to the front, and in the last four to somewhere in the middle. This means that a ligature can combine with the next character into a new ligature. In practice only the first variant is used and all of these work on pairs. There are very few fonts in the \TeX distributions that use the others. In general, the number of ligatures in traditional \TeX fonts is small, simply because the 256 slots we have there are easily exhausted. But it shows that in principle complex replacements are possible.

In OpenType fonts we also have ligatures but here they are part of several substitution features.

substitute	replace one character by one other character
alternate	replace one character by one taken from a set of alternatives
ligature	replace multiple characters by a single character
multiple	replace a single character by multiple characters

In addition to simple replacements like these there are also contextual lookups that look at (a range of) characters, and are optionally triggered by what comes before and what comes after and then perform one of the above substitutions. It's a bit off-topic but \TeX was written with Latin scripts in mind, so imagine what is needed to handle for instance Arabic with its much more complex ligatures. Fortunately hyphenation is not commonly used there so we don't need to deal with ligatures crossing the pre-, post- and replacement parts of a discretionary.

In the table above we have a ligature substitution feature but that name is misleading. For instance an '*ff*' ligature can indeed be a single glyph but quite often in OpenType it is implemented as two glyphs with some kerning, or with the second glyph being replaced by one that can blend with the first, which means that we end up with two glyphs with some negative kerning in between. The construction of ligatures can even be script- and language-dependent. But still this combination of single replacement

and positioning is called ``ligature'' and put under the `liga` feature. In `TeX` on the other hand ligatures are just multiple to single replacements.

As mentioned, the reason for this article is that there are users who fear that this ligature mechanism can be abused for malicious purposes. For instance, that we can replace one word by another word. A natural question for `TeX`ies is then: does OpenType make users more vulnerable than when they'd use traditional fonts? As we explained before, it will be clear that even with simple fonts one can do unexpected replacements but in both cases a multiple to single replacement is not what achieves this: *we need a multiple to multiple replacement*.

25.3 How it can be done in OpenType

Let's start with OpenType. Here we can do a multiple to multiple replacement using a chain substitution. We use the `ConTeXt` (font loader) `Lua` interface to define a new feature.

```
\startluacode
fonts.handlers.otf.addfeature {
    name      = "mytesta",
    type      = "chainsubstitution",
    lookups  = {
        {
            type = "multiple",
            data = {
                ["g"] = { "b", "a", "d" },
                ["b"] = { "g", "o", "o", "d" }
            },
        },
        {
            data = {
                rules = {
                    {
                        current = { { "g" }, { "o" }, { "o" }, { "d" } },
                        lookups = { 1 },
                    },
                    {
                        current = { { "b" }, { "a" }, { "d" } },
                        lookups = { 1 },
                    },
                },
            },
        },
    }
}
\stopluacode
```

Here **current** is a list of characters to match and every position can have multiple matches. When there is a match the lookup sequence points back to the lookups table where one of the mentioned substitutions can take place. This newly created feature can now be applied:

```
\definefontfeature[mytesta][default][mytesta=yes]
\definedfont[SerifBold*mytesta sa 1.2] good or bad
```

Here the default feature sets enable ligature building (**liga**), kerning (**kern**) and other features considered standard. We also default to so-called node mode because instead of using the **TeX** engine, we delegate feature handling to **Lua**. The output is:

bad or good

Indeed the words **good** and **bad** get swapped. We can add a **before** and **after** sequence too and you can imagine that when this logic is translated into a binary OpenType format a regular user will not know that this can happen. So, indeed a font designer or someone who can intercept and replace a font can hide replacements like this. If we replace **mytesta** by **liga** it becomes part of the regular everyday ligature feature and because that one is often enabled by default it could go unnoticed for a long time. Because OpenType fonts can be huge and have lots of rules like these this cheat can be hidden very well. To some extent **TeX** users are somewhat better protected because they can always turn features on and off, while in less configurable situations defaults are used.

So, to summarize, in OpenType we do this: as we run over the list, at every character we check for the sequences **good** and **bad**. When we have a match we replace the matched characters by something else. The next variant shows a bit more complex setup with more lookups and different substitutions.

```
\startluacode
fonts.handlers.otf.addfeature {
  name      = "mytex",
  type      = "chainsubstitution",
  lookups  = {
    {
      type = "substitution",
      data = {
        ["T"] = "t",
        ["E"] = "e",
        ["X"] = "x",
      },
    },
    {
      type = "multiple",
      data = {
        ["L"] = { "c", "o", "n", "t", "e", "x", "t" },
      }
    }
  }
}
```

```

        ["l"] = { "c", "o", "n", "t", "e", "x", "t" },
    },
},
},
data = {
    rules = {
    {
        current = {
            { "t", "T" }, { "e", "E" }, { "x", "X" }
        },
        lookups = { 1, 1, 1 },
    },
    {
        current = {
            { "l", "L" }, { "a", "A" },
            { "t", "T" }, { "e", "E" }, { "x", "X" }
        },
        lookups = { 2 },
    },
    },
},
},
},
\stopluacode

```

This might look a bit intimidating but it's just a way to enhance (or fix) a font and is in fact used in practice. This newly defined feature can now be applied:

```

\definefontfeature[mytex][default][mytex=yes]
\definedfont[SerifBold*mytex sa 1.2]
Using TeX (or TEX) logos, like LaTeX or (LaTEX), sometimes can look
a bit silly.

```

Using tex (or tex) logos, like context or (context), sometimes can look a bit silly.

Because we have control over the engine we can add more substitution features, not present in OpenType, because after all, control is what users love \TeX for.

25.4 How it can be done in traditional \TeX

In traditional \TeX we don't have such a repertoire of replacements available but we can use a different trick. Again we use Lua to emulate it.

```

\startluacode
local function initialize(tfmdata,value)
    if value then

```

```

local characters = tfmdata.characters
local y = utf.byte("y")
local e = utf.byte("e")
local s = utf.byte("s")
local n = utf.byte("n")
local o = utf.byte("o")
local cy = characters[y]
local ce = characters[e]
local cs = characters[s]
local cn = characters[n]
local co = characters[o]
characters[y].ligatures = {
  [e] = { char = 255 }
}
characters[255] = {
  ligatures = {
    [s] = { char = 254 }
  },
}
characters[n].ligatures = {
  [o] = { char = 253 }
}
characters[254] = {
  width    = cn.width + co.width,
  height   = math.max(cn.height, co.height),
  depth    = math.max(cn.depth, co.depth),
  commands = {
    { "char", n },
    { "char", o },
  },
}
characters[253] = {
  width    = cy.width + ce.width + cs.width,
  height   = math.max(cy.height, ce.height, cs.height),
  depth    = math.max(cy.depth, ce.depth, cs.depth),
  commands = {
    { "char", y },
    { "char", e },
    { "char", s },
  },
}
tfmdata.properties.hasligatures = true
tfmdata.properties.virtualized = true
end
end

```

```

local specification = {
  name      = "myhacka",
  description = "myhacka",
  manipulators = {
    base = initialize,
  }
}

fonts.handlers.otf.features.register(specification)
\stopluacode

```

Because in the end we need multiple characters we create a virtual font where a single slot (character) renders multiple characters. Of course one should also add proper kerning between the characters that make **no** and **yes**; copy the rightmost characters to the **kerns** table, and patch the ones that kern with the left character but we leave that as an exercise.

```

\definefontfeature[myhacka] [mode=base,myhacka=yes,liga=yes,kern=yes]
\definedfont[SerifBold*myhacka sa 1.2] yes or no

```

Where in the OpenType example we delegated processing to Lua, here we let the **TeX** engine do the job. This is called base mode. Node mode picks up the data it needs for ligaturing and kerning from the font via Lua; in base mode it is part of the font definition that gets passed to the engine (like the **ligatures** tables here).

no or yes

In the previous example we do this at runtime but if you use for instance **pdfTeX** you need to make a real virtual font file (a **vf** file alongside the **tfm** file). Here the **commands** entry provides the recipe for the virtual character where we draw from the same font.

So, to summarize this approach: we check each character for being the start of a ligature. When we have a match we replace the pair by a new character, in the case of **no** this is the final one, and in the case of **yes** an intermediate one that gets a follow-up check. Do you see the problem? If we enter this:

```

\definefontfeature[myhacka] [mode=base,myhacka=yes,liga=yes,kern=yes]
\definedfont[SerifBold*myhacka sa 1.2] yes or no or yep

```

no or yes or ^ap

We get an unexpected replacement, because character 255 is already taken. The solution is to add this:

```

characters[255] = {
  ligatures = {
    [s] = { char = 254 }
  },
}

```

```

width    = cy.width + ce.width,
height   = math.max(cy.height, ce.height),
depth    = math.max(cy.depth, ce.depth),
commands = {
  { "char", y },
  { "char", e },
},
}

\definefontfeature[myhackb] [mode=base,myhackb=yes,liga=yes,kern=yes]
\definedfont[SerifBold*myhackb sa 1.2] yes or no or yep

```

no or yes or yep

25.5 It's trickier than this

It's clear that where in OpenType one can use different means than ligatures, in traditional TeX where we only have ligatures as a tool we have to combine them with another (powerful) magic: virtual fonts. In the end there is not much to worry about: access to the font and the user's machine is needed in order to achieve this. In fact, the previous examples are not even right.

There is one aspect that we didn't mention yet that makes an OpenType font more suitable. If we only want to replace words we need to add a boundary check so let's give an example of that.

```

\startluacode
local punctuation = { { " ", ".", ",", ":" , ";" } }

fonts.handlers.otf.addfeature {
  name      = "myteste",
  type      = "chainsubstitution",
  lookups = {
    {
      type = "multiple",
      data = {
        ["g"] = { "b", "a", "d" },
        ["b"] = { "g", "o", "o", "d" }
      },
    },
  },
  data = {
    rules = {
      {
        before  = punctuation,
        current = { { "g" }, { "o" }, { "o" }, { "d" } },
      }
    }
  }
}

```

```

        after    = punctuation,
        lookups = { 1 },
    },
    {
        before   = punctuation,
        current  = { { "b" }, { "a" }, { "d" } },
        after    = punctuation,
        lookups = { 1 },
    },
},
},
},
}
\stopluacode

```

This can now be applied:

```
\definefontfeature[myteste][default][myteste=yes]
\definedfont[SerifBold*myteste sa 1.2] goody or badass and good or
bad
```

goody or badass and bad or good

Getting this to work in \TeX is hard because we run out of slots and because \TeX in base mode (traditional) has no space. And even in OpenType we probably need to do more to get it right all times: punctuation, digits, all kinds of quotes, etc. So with the given \TeX example we get even more unexpected changes because every work that contains the lookup replacement happens.

25.6 Unexpected ligatures in practice

Of course in practice fonts don't get hacked this way. But you can have surprises in some fonts. An example is Font Awesome Brands:

```
\definedfont[FontAwesome5Brands-Regular-400*default]
dropbox windows android apple linux wikipedia
❖ windows android apple linux wikipedia
```

Sometimes fonts have such goodies built in and as long as stylistic features are used for this it is just a handy feature (or you might call it a gimmick). Emoji fonts are good candidates for combining emoji into new ones, but changes in shapes (this could be driven by social or political influences) can also make for surprises. It demonstrates that it is always good to proofread what you typeset. The problem with such features is that when you decide to use another font, it is unlikely to work there so maybe it's best not to rely on it and use similar features that the macro package (might) provide instead.

A recent trend is to let monospaced fonts perform ligaturing, so we get for instance a ‘*less equal*’ symbol instead of a ‘*less*’ plus an ‘*equal*’. And it can be highly annoying if that can’t be turned off. Feature creep is probably more troublesome than an occasional nasty font, one that probably isn’t worth using anyway.

25.7 Some final words

The suggestion in the post that fonts are programs is not entirely true because we’re only dealing with lookups: it’s data that we’re talking about. That is not to say that fonts can’t become programs. As explained in Niklaus Wirth’s “*Algorithms + Data Structures = Programs*”, one needs the algorithms too and that’s what the OpenType composition code does, with specific features either supported or not by a user interface or hard-coded assumptions. And although the descriptions of OpenType have improved over time, occasionally one has to ponder how to deal with complex combinations. And as usual bugs can become features. In that sense, the limitations of the \TeX model are a protection against unexpected (but intended) side effects. If you nevertheless consider an OpenType font to be a program, you also should expect it to come with a manual explaining the reasoning behind the way it handles all this as well as showing the user what to expect, but I doubt that this is on the mind of a font developer.

So, while it is for instance no big deal in LuaTeX to add functions to the repertoire of features we can go further. At Bachotex 2024 I was told that there are (premature) ideas about adding Web Assembly to fonts, which for sure will give some users the creeps, but then we’re talking application-specific trickery which is outside our scope. The most recent color extensions are an example of undesired complexity and the question is if we want more of that. Maybe it makes sense on the web but on paper (and with \TeX) we can often achieve the same independent of fonts that then need to have these programs embedded in every instance (one can wonder how updating and bug fixes works out then). Once we lock our workflows into libraries that themselves have dependencies we also enter the realm of programs with potential back doors, leaks, insecurity, compatibility issues, etc. But that’s for another article. I only mention this in order to stress that the currently provided OpenType features are the least of our worries.

When you use a programming language like C or Pascal, the language \TeX was originally written in, there can be various results, most of which have little relation to each other. A word processor written in some language is not the same as a program that controls traffic lights. In addition to the core language, which provides data types, loops, conditionals, etc. there can be specific features like native support for strings, or a defined set of libraries to be provided, for instance for memory management or math. At some point it was fashion to add layers of abstraction to languages, for instance object oriented layers, or database oriented functionality. One can add to a language or use pre-processors for that. One can talk macros and templates.

If we look at \TeX we also have a core language but at the same time there is a rather large set of built in functionality. You can mix programming with content and the content gets processed into something that eventually might end up on paper or on a screen. The built in functionality deals with paragraphs of text, rendering math formulas, splitting of pages, making tables, and so on. It's actually an attractive mix.

When \TeX showed up the idea was that one writes a specific set of macros that help to structure and render the input. Although there is a basic set of plain \TeX macros, the look and feel got programmed per document. However, as not all authors have the skills to program or want to be bothered with that larger sets of macros popped up: macro packages. Some survived, some went extinct, but in the end we only find a few.

If you go back in time, the users of \TeX found each other in shared interest and a common need for resources. Of course there is the program, that needs to be available for the platform that one uses, then there was the backend driver needed for previewing or printing, and of course fonts and hyphenation patterns could be shared too. There was even some common ground for further development. At user group meetings it was not uncommon to hear about something new, successes and failures, usage patterns that can inspire etc.

But over the decades things changed. New engines showed up that added functionality. For instance LuaTeX brought an efficient mix of \TeX , Lua, and MetaPost, and LuaMetaTeX adds to the language as well as typesetting functionality. The first one resulted in the more hybrid ConTeXt variant MkIV , and the second in MkXL (aka LMTX as it pairs with the engine in a lean and mean installation). The later is what most users nowadays use. And the principles can be summarized as:

- We use the \TeX input method by default, which means commands that start with a backslash. As always with ConTeXt commands can have options: keywords (arrays), key/value pairs (hashes), or both. We do also accept xml, Lua (so called cld documents), some other formats and a mix of it.
- It being a major \TeX features, we support math in the usual way but we don't want users to tweak to much which is made possible by the enhanced engine.

- Solutions not available via the more high level user commands can be programmed using the regular \TeX language, Lua, or MetaPost or a combinations. There are of course helpers but we don't discourage going primitive.
- We consider \TeX as a niche product. For sure it excels at math but for many it is not beforehand a better solution than alternatives like word processors and web based rendering. It should never be enforced. Using $\text{Con\TeX{}t}$ is best a positive choice.
- It is users and usage that drives what features are added, extended, upgraded. As development and support is volunteer work there is also no or at least little (occasional projects) that binds it to big tech, publishers and hidden backend usage.

So how does $\text{Con\TeX{}t}$ then fit into the larger \TeX picture. As mentioned above in the beginning there was common ground but today one can wonder if that is still the case. With for instance \LaTeX moving to the web and competing with other large scale word processing, there is little in common. Also, we never had the same approach to styles (low level hacking) and definitely don't want to enforce some intermediate programming layer hides (and discourages) the use of the primitive language, although we do impose some restrictions with respect to overloading. We don't want to end up in competition with other systems out there either: one should be free to use whatever ones likes. There is no overlap in development and macro package development is rather isolated. You can see that reflected in use group meetings: they are small compared to the early days. The generic meetings bring little news and the more dedicated ones are (indeed) dedicated to specific solutions (macro packages). There is still some common ground to be found, like in distributions, specific resources like hyphenation patterns, and journals but even there macro packages go their own way. One has to look deep down to see to what extend they share concepts and usage patterns. Of course this is natural for a program that has evolved over four decades and where usage patterns also evolved.

So to come full circle, if we see \TeX as a programming language it is not much different from other languages. Of course its domain is typesetting but in what way the results are achieved can differ a lot. When someone says to be using \TeX that doesn't tell much more than saying that one programs in Pascal, although it is an indication of how typesetting solutions can be reached.

In a nostalgic mood I took a look at the rather large archive of Modula files I wrote in the second half of the 80's and first half of the 90's. Among the dozens of programs are a few that sort of translate to what we have today. We're talking msdos here using a common text based windowing system that I wrote for vax mainframes and ported to the personal computer.

There is an editor that could handle relatively large files, supported a project structure, had syntax highlighting and did real time spell checking. I just translated that to SciTE extensions and for instance the project structure (including fast opening of files) is not much different from what we have now. The syntax highlighting evolved to a mixture

of languages but still determines how I look at files and thereby how ConTeXt is programmed.

But before we actually used TeX that same editor was used for just ascii based rendering so in addition we then had a converter that turned simple tagged text into something for a printer. It actually is not that different from these markdown ways of coding. We also could split of pages and move reserved areas (floats) around. Of course all that was left behind when we moved on to TeX. It is kind of curious that much of this actually was written for terminals connected to a VAX (at the university) where we printed out on these fast daisy wheel line printers. Only decades later we figured that already at that time there could have been TeX or precursors running on those machines, probably of little use because there one needs an (expensive) phototypesetter and some workstation for previewing.

Yet another program was the one that we used for making computer assisted learning programs, think of presenting text, questions, feedback, video and audio. Equipment was controlled via serial connections connected to for instance random access video tape players. We made a small language for that with regular programming features enhanced with some specific for learning environments. Interesting is that it involved parsers, scripts in that language, efficient storage, and bytecode ran through what I later learned was called a virtual machine. I guess that around that time, when we were not connected to the (not yet existing) web, many people were inventing the same wheels. As a side note: there was never a real market in a small country like ours so this was typically only applied in projects where it somehow fit in, as goodie. Of course all this is now done in browsers and JavaScript, but I guess the experiences somehow translated into interfaces mixed into ConTeXt. I never felt the need to abstract TeX away behind layers like the one mentioned here because nowadays we have these interpreted scripting languages.

Then there were graphics. At some point we were involved in making a heavy duty (3D) milling machine using simple robust of the shelf components. The friend we did this with needed it for some production process. We controlled motors with brakes and couplers and feedback happened with optical feedback. All was programmed in Modula and controlled via an io board. Here again we made some simple language for describing shapes but also handled output from drawing programs (exporting to hpgl). Looking backward it was a lot of fun and there was nothing on the market that could do the same for that price but the (more than square meter devices) were mostly used for milling displays and components for an automated film development machine. Here we can wonder how that translated to today's ConTeXt. Using MetaPost for this is not an option although with the library we have we could actually cook up some backend. Also, you can now use micro controllers that cost a few euros and stepper motors than come cheap. So, looking back it's probably mostly the learning experience that translates. And the recent serial ConTeXt signal gadget brings us back to connected hardware.

There are more programs, some concern databases, others concern for instance distributing structured handbooks (on floppies, meant for portable simple laptops). One

easily forgets what has been done. The same is true for manuals: we can find plenty of early ConTeXt manuals dating from the time it evolved. Most never made it into the public. Much also had to be done on computers from the i386 times. And it definitely was from before we were connected to the internet. But in the end, with intermediate steps using Perl and Ruby for scripting and run management, we eventually arrived at this TeX, Lua and MetaPost mix, now with even a bit of connected hardware to watch the process. There is no need to look back apart from realizing that we didn't arrive here by accident.

So to wrap this up: at some point we arrived at using TeX combined with various programs we made that themselves evolved but eventually all those became obsolete and the Lua plus TeX plus MetaPost approach was the winner, with an occasional resemblance of what came before. First we operated in isolation, then entered the TeX community that at that time was also exploring several solutions for similar problems and adapting itself to what it was confronted with. But in the end, just as we forgot the past, that TeX related past also starts fading away: the current TeX solutions evolved independent and share little. But what remains is that they are built upon Don Knuth's masterpiece and that we share forever.

comment: this text is not yet corrected for errors

27.1 Introduction

I write this 20 years after we started the LuaTeX project, which was 2005. At that time I used pdfTeX, for obvious reasons: it was a mature engine, had a pdf backend built in, came with nice features like expansion, had a few extensions at the primitive level, and I was involved in the development, testing, discussions etc. I was quite happy with it and ConTeXt always supported the latest greatest smoothly. And after decades MkII still works well with the current pdfTeX, which has evolved a little but we never followed up on that.

During the 20 years we reported frequently about how we moved on to LuaTeX, which got frozen after some 15 years but got a follow up in LuaMetaTeX. Among the reasons mentioned for starting the project were the potential to extend the inner working of a TeX engine via a Lua interface, the fact that the world moved on to Unicode and Open-Type fonts becoming more popular and widespread. We also had to support right-to-left typesetting in a more comfortable way which is why the Oriental TeX project became instrumental in the first five years.

In the next sections I will reflect a bit on how the code base evolved and how that fits into how the average TeX code base looks like. The bottom line is that we moved from literate to ‘illiterate’ programming.

27.2 Woven webs

If you look at the TeX live codebase for the familiar engines of course your first try to find the original. In the October 2025 snapshot we find:

```
1,031,999 tex.web
180,075 tex.ch
```

Here the big file defines the program and the file is big because it has program code as well as documentation. If you open the file you’ll notice that it’s quite dense. If spacing had been more liberal and the various directives more verbose the file would be way bigger but think for a moment about those times: the first 10 MB hard disks in personal computers that showed up years after TeX was written were larger than a modern mini pc and weighted more than a decent laptop. You not only have that 1 MB text file, you also need to compile it, so you need a software infrastructure, you need to test the binary, which also takes space, save a memory dump, collect some fonts, generate bitmaps, etc. It’s a miracle all worked out so well. Oh, and we didn’t mention that in those times printing hundreds of pages in a decent resolution took some time too.

The small file is what is called a change file. From the big `.web` file we can ‘weave’ documentation (like the \TeX the program book) and ‘tangle’ a Pascal program. The change file is applied in the process and it patches the to be woven web. Such a change file typically deals with system dependent issues but can also add or upgrade features. If you keep in mind that the first version of \TeX was configured for 32 bit memory words, it will not come as a surprise that the first change was going 64 bit so that the engine for instance could handle more so called nodes.⁴¹

An example of change file that adds functionality is e- \TeX , one of the first extension projects:⁴²

`225,207 etex.ch`

That one is about a quarter of the original. There is no `.web` file because it patches the original \TeX file. This is not the case for pdf \TeX :

`1,577,910 pdftex.web`
`16,272 pdftex.ch`

Here we see a copy of `tex.web` with adaptations but still there is some change file. It deals with the fact that the `tex.ch` file is also applied as well as some m \TeX functionality. Keep in mind that we operate in an infrastructure where code is shared, especially interfaces to the file system and environment (`web2c` & `kpse`). The pdf \TeX program has a backend and interfaces to various graphic libraries and all those are done in C. So, in addition to these `.web` and `.ch` files we also find `.h` and `.c` files. We have a mix if literate and less literate code.

We see the same with X \TeX , where due to less added functionality (code and documentation) the web file has grown less but that is compensated by more of the new front-end functionality done in C. Also, keep in mind that this program delegates the pdf backend to dvipdfmx by piping out dvi.

`1,378,780 xetex.web`
`31,800 xetex.ch`

An interesting case is Aleph, the production version of Omega. At \TeX meetings this multi-directional, large font variant was often presented in a way that made it sound like it was not coded like \TeX due to various limitations, e.g. hardware. It would even be in a transition to C. However, when you look at the code base in $\text{\TeX} \text{Live}$ you’ll notice that it is all change files and a few C files. It actually stays within the concept of patching the original and where in one way it does away with limitations that other engines work around, it otherwise is kind of conservative. Of course it might relate to the fact that the added functionality is in transforming input and feeding the backend. For instance,

⁴¹ These are the times of big and huge \TeX ’s and hooking into memory extenders on 640K personal computers. The times when processing pages could make for a coffee break.

⁴² I first heard about e- \TeX and Omega at the first \TeX conference I attended, 1995, Arnhem NL. I was a bit amazed that there was this competition going on.

the directionality doesn't really need many changes in core code. The 35 change files add up to 813,829 bytes and files range from 124 to 174,831 bytes.

We end with MetaPost. Here we have no `.web` files but `.w` files because this program is not written in Pascal but in C. There are no change files in its directory. Some functionality is delegated to specific files, like the number systems that are supported. The binary, decimal and double variants depend on other code, so in the end we do use additional header and code files. The optional `png` and `svg` backends also add (quite huge) dependencies.

```
1,213,713 mp.w
  63,738 mpmath.w
  68,096 mpmathbinary.w
  73,831 mpmathdecimal.w
  52,224 mpmathdouble.w
  45,035 mpost.w
  12,369 mpstrings.w
```

Here the big one is actually the library and the small `mpost.w` implements the command line program.

27.3 Changing the web

In the previous lists of programs we miss `LuaTeX`, so let's look at that now. Because we started from `pdfTeX` the first version was just a copy and therefore quite similar. The `Lua` interface gave access to a few internals and we could write back to the input. Of course by adding `Lua` some `C` files were added. Apart from some extra primitives the `TeX` system was just the same.

However, as we decided to go forward and really open up the internals it became clear that the single file approach would not accommodate that well. But, with future `Lua` interfaces being `C` driven and with upcoming demands with respect to fonts in mind, it made sense to move away from `Pascal`. Keep in mind that for instance for `oriental TeX` we had to load relatively large fonts (for which we initially took code from `FontForge`) and process the so called node lists that makes boxes and paragraphs. We didn't want to go the `XeTeX` route by depending on operating system features and/or specific libraries that could change (and in the end did change).⁴³ For us, `Lua` was the way to go: after all, what is a standard (fonts, pdf, graphics) if they lock into specific libraries; we wanted a programmable extensions system, pretty much like `TeX` itself.

So, to summarize a few years into a few sentences: `Taco` and `Hartmut` split up the code base, moved from `Pascal` web (`.web`) to `C` (but kept the comments) and eventually ended up with `C` web (`.w`) files. The idea was even to write new code as a web but in

⁴³ We do support `XeTeX` in `ConTeXt MkII` but never used its features beyond loading fonts because they didn't match our approach. So, I can't really comment on how it works internally.

practice that was just too much work and never really happened. It anyway meant that we no longer used the web2c route but the cweb route.

This process actually took years because at the same time we also added functionality, improved some existing mechanisms (like more efficient expansion), did some cleanup, etc. We tried out ideas, implemented solutions in Lua, adapted as we went and had to document as well as report on it: see ConTeXt documentation, TugBoat, Maps, etc. It was very convenient that we let ConTeXt MkII evolve into MkIV and that ConTeXt users just joined in testing. Basically LuaTeX could always be used in real production flows, as long as one kept ConTeXt and the engine in sync.

When we decided to more or less freeze the engine, around version 1 in 2016, we had a mixed code base, which was a good moment because Taco no longer could allocate time; keep in mind that this is all unpaid volunteer work. After that I started to add pending functionality and some more interfaces, while Luigi took care of integration in TeXLive. It kept us busy for a few years. In the process we change the `.w` files to just `.h` and `.c` files and keep the comments as such. We didn't expect anyone to ever generate a pdf from the source anyway, also because it never could match the original Knuth quality. One needs discipline for reaching that level and the combined workload (TeX, MetaPost, macro package, documenting, usage) simply didn't provide room for that. Also, editors with syntax highlighting, cross file searching, grepping, high resolution windowed display technology, and whatever, had evolved and looking at the documented source was as good as looking at a (to be generated) pdf on screen.

There was an exception to the rule: MetaPost. That big file was already in documented cweb and splitting it up is a major effort. It did mean that we had a dependency on a converter but we'll see how that worked out.

27.4 Becoming illiterate

We've now arrived at LuaMetaTeX. When the more definitive code freeze of LuaTeX was decided in 2018, the development moved to a new variant. At that time the code base can be summarized as follows. There is a set of files that can be divided in those implementing the TeX engine and user interface (interpreter), Lua libraries that are kind of stand alone, Lua libraries that interface with the engine, libraries that give access to fonts or generate pdf, and of course the MetaPost library. Even libraries dealing with bitmap images could be replaced by Lua solutions. The files are stored in a TeXLive compliant structure, and use its build infrastructure. It's a bit messy mix. An optional library interface makes it possible to add specific features via Lua but ConTeXt will never depend on that.

Among the decisions made when I started LuaMetaTeX was that the font loader code could go away, because after all in ConTeXt we already did all of this in Lua. The same was true for the backend: we experimented that in LuaTeX in MkIV (but later removed that option there). The decision also made for a new version of ConTeXt: MkXL or, as distribution LMTX, because at that time it became clear that we had to distribute both at the same time.

I was surprised to find out that in the end compilation was not that hard, not even with the native windows compiler.⁴⁴ It made for a way more convenient development cycle but there was one dependency left: MetaPost was still webbed. In order not to be dependent on the web2c toolchain I wrote a Lua script that did the same and also had the possibility for plugging in Lua code to fix certain issues. A side effect was that I got better formatting of the C (which presumably was not meant to be looked at, just compiled) and a few compiler warnings related to indentation. When I really started to extend MetaPost the `.c` file was used and back ported to `.w` so after a few years, sentiments set aside, the cweb was dropped from the source tree. It triggered another round of cleanups.

So, is there anything literate left? The answer is “*partial*”. The various comments are spread over the sources as C comments. They come from original \TeX , $\text{pdf}\text{\TeX}$, $\text{Lua}\text{\TeX}$ and of course from the author. The original comments do not always apply but there is also a history to be told and the general ideas and approaches of \TeX don’t change. We deliberately always tried to stay close to the original when possible.

I have it on my todo list to clean it up a bit more. The challenge is to keep the explanations, even if they no longer apply that way. History also matters, and it’s not like many people read it. A better source of explaining what the engine does are the articles, development documents, low level manuals and $\text{Con}\text{\TeX}t$ code base. It’s more important to write for users than for whoever looks at the source.

What I didn’t stress enough is that the code base is independent. There is some third party module code, like (of course) Lua, memory management, compression, pdf parsing, specialized graphic stuff related to MetaPost but all of it is in the tree and distributed with $\text{Con}\text{\TeX}t$. Updates never happen automatically, changes are checked (diffed) first. When a user installs cmake and a compilers building a binary is easy. The codebase can quite well be explored from Visual Studio, VSCode or likewise. (But rest assured: it’s free of AI interference.)

27.5 Conclusion

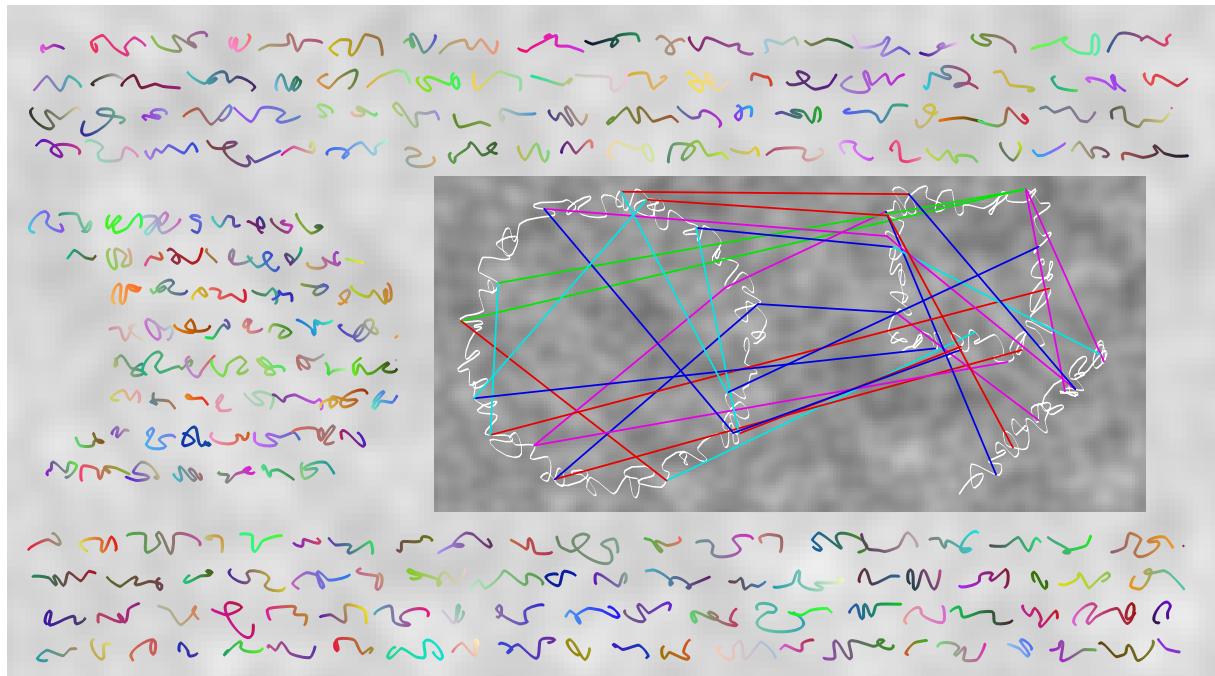
The $\text{Lua}\text{\TeX}$ and $\text{LuaMeta}\text{\TeX}$ code bases differ from other \TeX engines. Where $\text{Lua}\text{\TeX}$ integrates in the $\text{\TeX}Live$ code base, $\text{LuaMeta}\text{\TeX}$ is independent. Both don’t use the web main file plus change files approach, although $\text{Lua}\text{\TeX}$ does so for `mplib`, which actually then has its own source tree. So, both don’t really qualify any longer as literate programs, but they do have plenty traces of the idea that the source should explain a bit what happens.

⁴⁴ When working on $\text{Lua}\text{\TeX}$, I cross compile on the Linux subsystem for Windows using a build script that Luigi made. The files are organized in the $\text{\TeX}Live$ way, and the build script fits into that. Cross compiling $\text{LuaMeta}\text{\TeX}$ is faster and for various platforms we use Mojca’s $\text{Con}\text{\TeX}t$ garden build infrastructure. Installing $\text{Con}\text{\TeX}t$ and $\text{LuaMeta}\text{\TeX}$ is done by the engine itself as it can act as Lua interpreter and therefore download and install files. It is a self-contained setup.

I never realized that LuaTeX is an outlier because it doesn't have this change file model. The objectives of the project also made us change for instance the size of nodes and font resources which means that there are lots of accumulated changes. In an other musing I talk a bit about the way character nodes are managed in original TeX and derivates. The change file model is a sort of insurance that the related optimizations are kept but also might limit the solution space. So it's a blessing on the one hand and maybe a curse on the others. I can safely say that we would not have gone as far as we went if we'd locked into the approach that the other engines took.

So, kudos to Hartmut Henkel and Taco Hoekwater for having the courage to divert soon after we started and of course thanks to Don Knuth for providing the well known and very literate documented source to do all this, the ultimate source I still go back to when in doubt.

In a ConTeXt development team discussion, Keith McKay offered to write some documentation and we all agreed to adopt his more futuristic style of literate programming. Here is an example:



After all, there is no need to stick to TeX alone and we have to make sure that only humans can get the picture what we want to achieve: it's all about beauty, not about turning what we do in artificial slop.

In these days of running into so called '*fake news*' it is no surprise that we also can run into fake quality typesetting. Of course, as with news, it is can be hard to figure out what actually is real or fake, so with typesetting we also need to be careful in what we claim. For that reason I will not mention names of programs, nor refer to claims made about how well programs behave, even when they compare themselves to for instance \TeX , in which perspective one can read strange assumptions with respect to what \TeX can do or not, nor will I comment on claims made by designers of journals who brag about the improved quality due to choice of fonts, applied trickery, design choices and what more. I just stick to some abstract visual examples that might help the reader to identify false claims, especially when to have to defend their use of \TeX and friends. What follows below is about awareness. Trust your eyes. If looking at some text gives so a feeling that something is off, wonder why. If it looks kind of weird, find a reason.

We start with the challenge. When typesetting a paragraph of text, a typesetting engine has to make choices. When words can be hyphenated, a potentially overfull line can be prevented by breaking a word. When spaces have stretch or shrink, that can be used to make better decisions. In the case of \TeX a first look at the paragraph might result in breaking without hyphenation. If the result is too bad, a second attempt will be made, this time with hyphenation. If we're still not okay, the criteria can be more liberal and we can permit additional stretch between (either or not hyphenated) words. When set up right, the final result will not have any overfull lines but spaces can be excessive in some lines. However, in practice that seldom happens. The engine looks at the whole paragraph, builds a solution tree and chooses the best possible outcome. That of course comes at the cost of runtime but nowadays that can be neglected.

Starting with pdf\TeX some more wiggle room became available: character expansion. This is a variant on so called hz optimization, a trick inspired by Herman Zapf, loosely based on Gutenbergs "*choose a wider or narrower variant of a shape*". In this case we permit a glyph to scale in the horizontal direction and we can use different scales per glyph. Of course that means that one has to set up fonts properly and the engine has to work harder because now it has to take potential stretch or shrink into account as well. The more constraints we have the more runtime is needed.

In Lua\TeX this mechanism has been optimized so there the impact is less, although more complex feature processing obscures that, which means that in the end this engine is more demanding when it comes to typesetting a paragraph. In LuaMeta\TeX we go even further by permitting more than the usual up to three passes: we can follow different scenarios to get to an optimal result. In practice the additional runtime for e.g. up to six passes can be neglected.

So, when we look at the result, and assuming that we have expansion applied, we can end up with a paragraph that have lines where no expansion took place, or lines where

glyphs became a bit narrower or wider, depending on what was needed to make lines visually compatible. And that last fact is what we need to stress: the \TeX engine tries to make neighboring lines look compatible! Again, in LuaMetaTeX we can do that for three, five, seven etc lines, where the distance between lines is taken into account, so we can be more tolerant the further away we are. And yes, that comes at a price in runtime, but as said, in practice one will be fine: runtime matters but getting a performance of 50 pages per second on my 2018 laptop for moderate complex documents, where typesetting the paragraph only accounts for a small portion of the process, I guess we're good. And it can only get better with more modern machines.

All that said, we now come to the point where we can look at what one can encounter in documents. For this we visualize two processes that we see applied in other applications. The first is inter-character kerning, the second is horizontal scaling. When we talk \TeX , the first method can be applied but that is normally only done in special situations, like some title. If it is needed for a regular text run one should make a better choice of fonts. The second process can be part of \TeX 's typesetting but as discussed before, it happens as part of decision making and is granular with respect to glyphs. It is not applied after the lines have been broken, it is applied as we decide to break. It is not a last resort, it is part of the decision making.

In order not to fake news, I will not make any assumption about how other applications construct a paragraph. I just don't know any other application than by name: I never use a word processor, page construction program, or any batch processing system other than \TeX . I just have no reason and application and am not even interested on how things are done elsewhere. When we develop solutions the driving force is not "*Can you do this or that as done in program such or so?*" but more "*We have this problem, can we handle that?*". And of course we just start and follow up on where \TeX ended up decades ago. It's all about having fun solving challenges.⁴⁵

But, if you look at some typeset page, you might wonder how it is done. For instance, when you see excessive horizontal stretch or space between characters but also notice that there are better ways to hyphenate, you can be sure that \TeX is not used. First of all, it does not expand kerns unless told so and even then does that within limits that go unnoticed. When glyphs get scaled normally, a macro package will put constraints on that, although I admit that I've seen badly done \TeX documents too.

So, when you wonder about quality of what you see, maybe it is due to the fact that the program used went on line by line, not much helped by proper hyphenation, and eventually started to apply extra kerning between glyphs either or not combined with stretch or shrink. The main objective likely was to get rid of too wide spacing and sacrifice quality of the text over that. It's like hammering screws into wood instead of using a screwdriver: it might work out but looks bad. In such cases just ask yourself: was the whole paragraph taken into account, was hyphenation properly applied, did

⁴⁵ Quite often we have applied \TeX in situations where other solutions failed. Manipulating and cleaning up bad and unstructured input was then part of the game. The fact that we could typeset decent results mattered less, it only made the work more rewarding.

anyone actually look at the result? The latter is actually a pretty valid question. If you consider workflows it can be that for instance in journal a lot of time went into the peer review and editing, then out-sourcing all to some typesetting company, assuming that all is okay and done to specs and finally goes to print. The programs used are supposed to do the best they can and be properly applied. The whole process is managed in steps, and following these steps is the assumed guarantee for quality. We can see results in scientific journals and books that really make you wonder if someone even look at the typeset results. Compare that to children books, where it is actually hard to find bad examples.

We can visualize the methods. We use simple sequences of glyphs and apply stretch or shrink as well as add some proportion of the width as kerns between. We use word of 9, 10 or 11 characters. These are not real text examples so we let the text overflow the line width for the sake of showing what happens.

We start with compensating spaces by making the words wider. The first line is what comes out naturally. The second line applies a bit of stretch. The next pairs of lines show what happens when we add kerns, with or without stretch. The spaces become less but the picture worsens.

In the margin we see three numbers: the width, kern and horizontal scale.

1.00 00 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 00 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 12 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 12 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 10 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 10 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 08 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 08 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 00 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 00 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 12 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 12 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 10 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 10 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 08 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 08 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 00 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 00 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 12 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 12 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 10 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 10 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 08 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 08 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB

The next examples show what shrink does. Maybe it is good to know that \TeX , when it comes to shrinking spaces never goes below zero but when stretch is permitted it can go beyond what is specified, the **plus** value is used relative to other stretch applied.

1.00 00 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 00 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 12 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 12 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 10 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 10 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 08 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 08 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 00 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 00 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 12 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 12 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 10 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 10 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 08 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 08 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 00 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 00 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 12 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 12 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 10 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 10 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 08 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 08 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 00 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 00 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 12 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 12 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 10 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 10 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
1.00 08 1000	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB
0.90 08 0900	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB	BBBBBBBBBBB

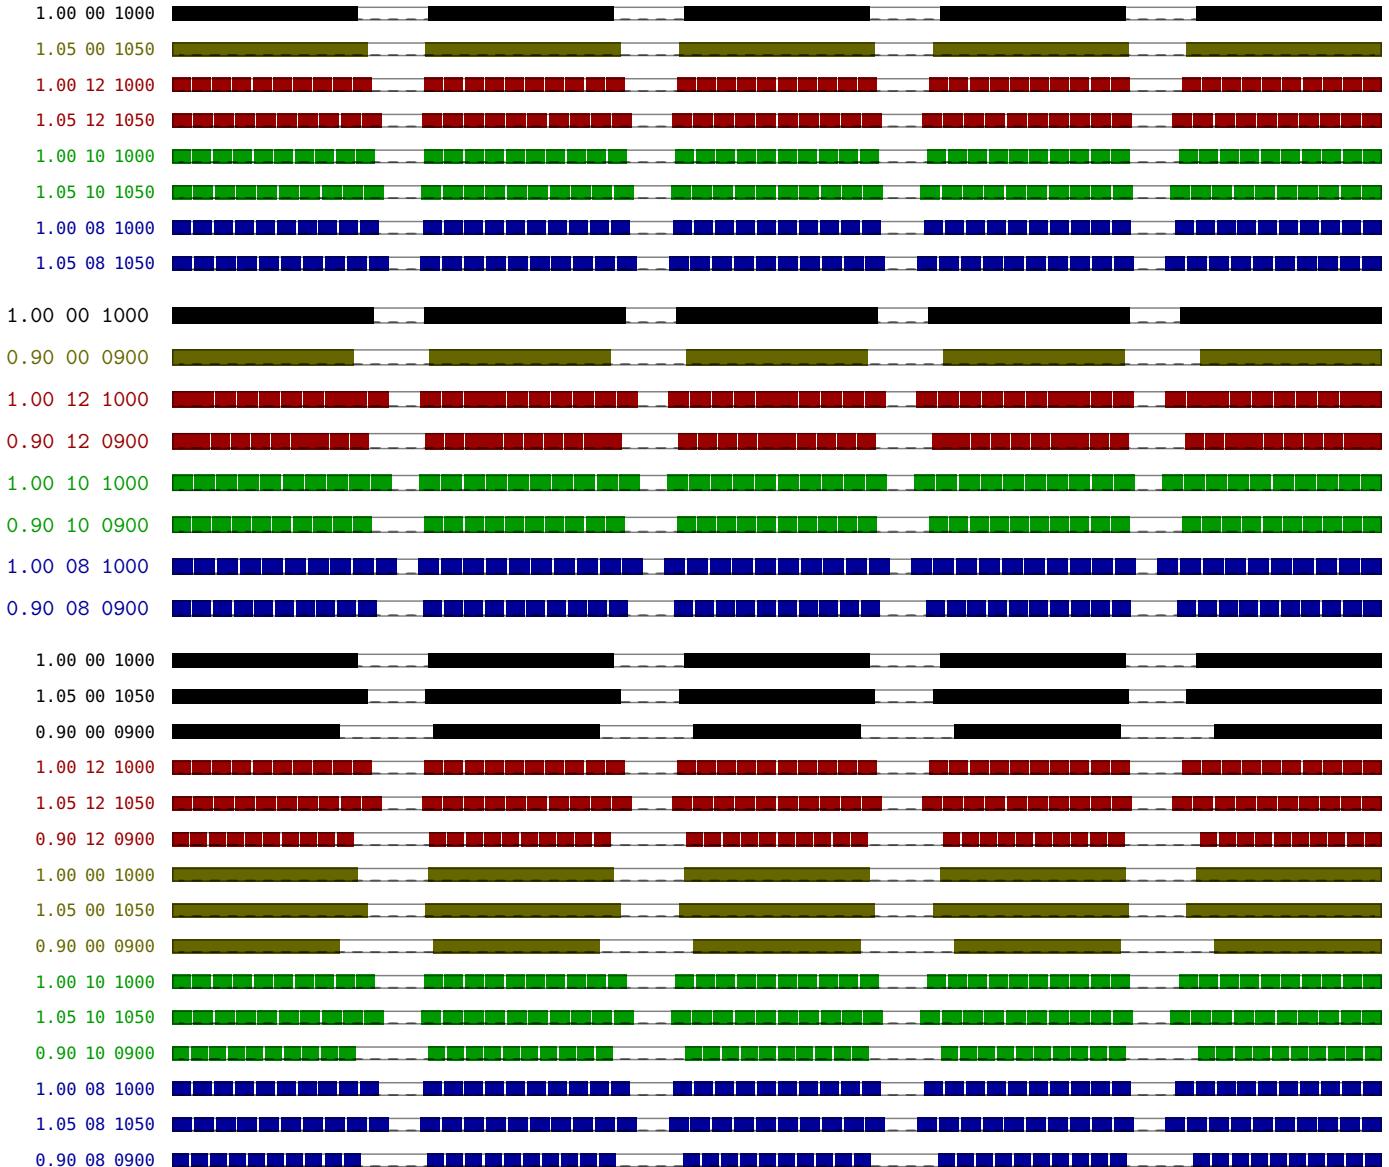
Last we show triplets of lines and those extreme differences is what one can observe when an application doesn't consider visual compatibility and permits such extremes. And yes, you can run into these cases in journals, newspaper and books. We only show one words size.

1.00 00 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 00 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 00 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 12 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 12 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 12 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 00 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 00 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 00 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 10 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 10 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 10 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.00 08 1000	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 08 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
0.90 08 0900	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB

Our main objective is to make you aware of these two tricks being applied by applications. I want to stress again that in \TeX this kerning doesn't happen, although it is able

to widen inter-glyph font kerns as defined by the kern features. One pitfall with adapting these is that it assumes that scaling doesn't impact for instance ligature building or accent anchoring.

We can also use rules instead of glyphs to demonstrate this, so we end with a few examples of that.



It can get real bad, and don't pretend that you never saw such extreme rendering, after all with a \TeX trained eye you can't avoid noticing.

1.05 08 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.25 08 1250	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.50 08 1500	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 05 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.25 05 1250	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.50 05 1500	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.05 03 1050	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.25 03 1250	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB
1.50 03 1500	BBBBBBBBBB	BBBBBBBBBB	BBBBBBBBBB

So how does TeX keeps up with the mainstream typesetting machineries? Here we have a torture test, a quote from Donald Knuth in Digital Typography.

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that “*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*”) Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j > 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that $(Ay + B)/(Cy + D)$ lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid’s algorithm, and terminates.

Here we show the bounding boxes of glyphs and the applied font kerns. You might need to zoom in to see them.

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that “*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*”) Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j \geq 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that $(Ay + B)/(Cy + D)$ lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid’s algorithm, and terminates.

Below we show the same text with expansion applied. You can see where it gets applied.

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that “*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*”) Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j > 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that $(Ay + B)/(Cy + D)$ lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately

switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid's algorithm, and terminates.

In LuaMetaTeX we actually also can apply expansion in math, which is shown next:

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that "*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*") Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j > 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that $(Ay + B)/(Cy + D)$ lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid's algorithm, and terminates.

Although it might not be visible in the next example we also use a more detailed paragraph pass setup, we use `\setupalign[mathbookpasses, hz]` here. The more advanced passes setup actually means that we only apply expansion when other solutions are worse, so we might not get it at all, but we leave it to your eyes to decide if that happens.

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that "*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*") Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j > 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that $(Ay + B)/(Cy + D)$ lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid's algorithm, and terminates.

Here we also enable expansion in math, so does it kick in here or not?

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that "*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*") Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j > 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that

$(Ay + B)/(Cy + D)$ lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid's algorithm, and terminates.

When we half the width we still get a reasonable result. If you still think in expansion only, realize that typesetting math involves more, especially when you don't want to break in certain places. We can also enable features that control how many hyphens we want in a row, how the last line should look, what penalties kick in, if we discourage similar words at the end of lines, or short words, or . . . In that respect often spacing is the least of our worries. In narrow paragraphs the likelihood of expansion kicking in increases.

15. (This procedure maintains four integers (A, B, C, D) with the invariant meaning that “*our remaining job is to output the continued fraction for $(Ay + B)/(Cy + D)$, where y is the input yet to come.*”) Initially set $j \leftarrow k \leftarrow 0$, $(A, B, C, D) \leftarrow (a, b, c, d)$; then input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$, one or more times until $C + D$ has the same sign as C . (When $j > 1$ and the input has not terminated, we know that $1 < y < \infty$; and when $C + D$ has the same sign as C we know therefore that $(Ay + B)/(Cy + D)$

lies between $(A + B)/(C + D)$ and A/C .) Now comes the general step: If no integer lies strictly between $(A + B)/(C + D)$ and A/C , output $X_k \leftarrow \lfloor A/C \rfloor$, and set $(A, B, C, D) \leftarrow (C, D, A - X_k C, B - X_k D)$, $k \leftarrow k + 1$; otherwise input x_j and set $(A, B, C, D) \leftarrow (Ax_j + B, A, Cx_j + D, C)$, $j \leftarrow j + 1$. The general step is repeated ad infinitum. However, if at any time the *final* x_j is input, the algorithm immediately switches gears: It outputs the continued fraction for $(Ax_j + B)/(Cx_j + D)$, using Euclid's algorithm, and terminates.

Our experiences with moving on with engine development, improving rendering of math, and typesetting in \TeX in general has shown that publishers seem to care less about these matters than one would imagine. At least, no matter what we publish about it, we hear little to nothing from those using \TeX who deal with their typesetting. So as \TeX ies we just do what we like best, keep a low profile, and focus on our own documents, which often happens to be in an educational settings where hopefully students still appreciate the look and feel and don't want to get distracted by bad typesetting.

Another trick that \TeX ies like to apply is protrusion. If applied, that must be done consistently which actually adds a constraint. I always wonder if that then makes for better results; after all the solution space doesn't become much larger. But here a typical \TeX user syndrome kicks in: because they know that \TeX is good at these things, they also assume that the result is good. However, the Internet demonstrates that plenty of users don't notice overfull boxes, sub-optimal math rendering and other artifacts. In that respect it is no surprise that journals that these users read also get a pass.

In ConTeXt (using LuaMetaTeX) you have a lot of control over the typesetting. For instance you can locally influence hyphenation, enable or disable font kerning and ligature building between (before and/or after) characters, influence scaling and what more. In practice users will not do that and just rely on the par builder doing its work, trusting fonts to be okay. It also defeats a bit the fact that users want this to be done automatic.

Given the above, we have to come back to the occasional "*I use this because it beats that.*" argument when it comes to typesetting systems. First of all, just use what you like. If seeing a backslash in a source annoys you, find a system with a different escape symbol. If you think that minimized encoding combined with tricky parsing solves your problem, go ahead. If performance matter, hunt for faster programs, but realize that, as the problem doesn't change, future versions of that solution might run slower when they do more. But all that doesn't change how you can look at the result because that is what matters. The TeX ecosystem has set some standards and inspired some developments but in many aspects remains a sort of reference.

- How consistent is vertical spacing: lines and between lines, around section titles, before and after paragraphs, around images, etc.
- Is horizontal spacing uniform: between words, around table cells, are proper font kerns applied between glyphs, is there additional kerning between specific glyphs (e.g. in titling) and is it done right, etc.
- Is math properly and consistently spaced, scaled, stacked, aligned, broken across lines, etc. As with text rendering, once you know what to look at you can't unsee it.
- Can you identify annoying interference, often visual by inconsistencies all over the place, weird artifacts, sloppy checking of overfull boxes.
- Are images consistently scaled, especially when they have embedded text, do fonts in them match the document, is bleeding intentional or just an oversight.
- Is the text properly aligned, has hyphenation been applied when possible to make for better solutions, is expansion used within reasonable bounds (read: close to invisible). It's easy to go overboard here and it can become a reason for not reading and/or buying books.
- Did the user constraint excessive font abuse, are features properly applied to fonts, do font sizes make sense. Does the font make reading a pleasure.

We can come up with more but you get the idea. Just look and be honest. A problem with TeX usage is that one can tweak a lot and thereby also mess up a lot. Bad spacing in text and math and bad font usage can be seen all over the Internet. The system can do a lot of good, and users (or styles) a lot of bad. The fact that TeX is used is no excuse for not looking carefully at the result, nor being critical. But after all these years I'm convinced that plenty users don't really care: they key in, render and never look back; they just assume it's okay. This is quite understandable because not every (maybe even

forced to use it) user cares about, is interested in, or even aware, of typesetting at all. However, that is not our intended audience but if you came reading this far, you might be.

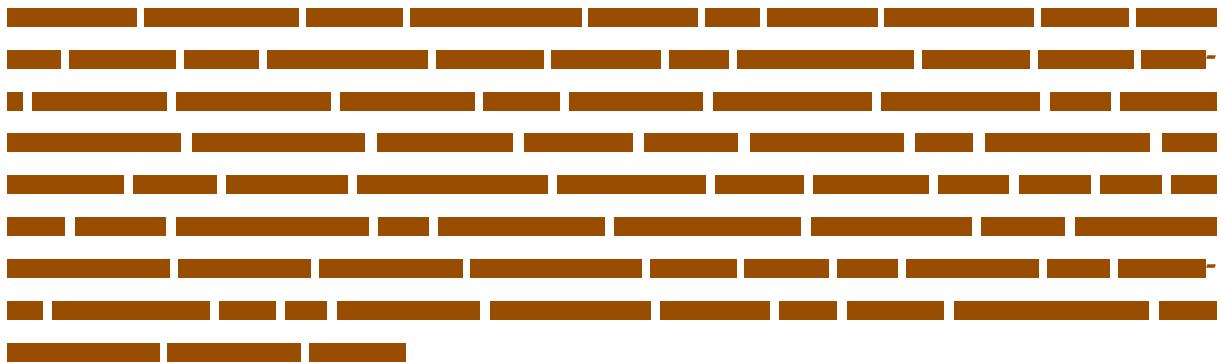
In case you want to know how a sample line is typeset, here is some code. We use EBGaramond as sample font, have a 12pt bodyfont, assume 65 characters per line and therefore have a fake width of 7.00372pt. This wrap-up is not a tutorial, it just wants to make the reader aware of what goes on and what to pay attention to

```
\starttexdefinition Fake #1#2#3#4
  \dorecurse {#1} {
    \ifconditional \FakeRules
      \vrule
        height 1\exheight
        depth .1\exheight
        width #2\FakeWidth
      \relax
    \else
      \glyph
        scale \FakeScale
        xscale #4
        \FakeCharacter
      \relax
    \fi
    \ifcase#3\else
      \kern{\FakeWidth/#3}\relax
    \fi
  }
\stoptexdefinition

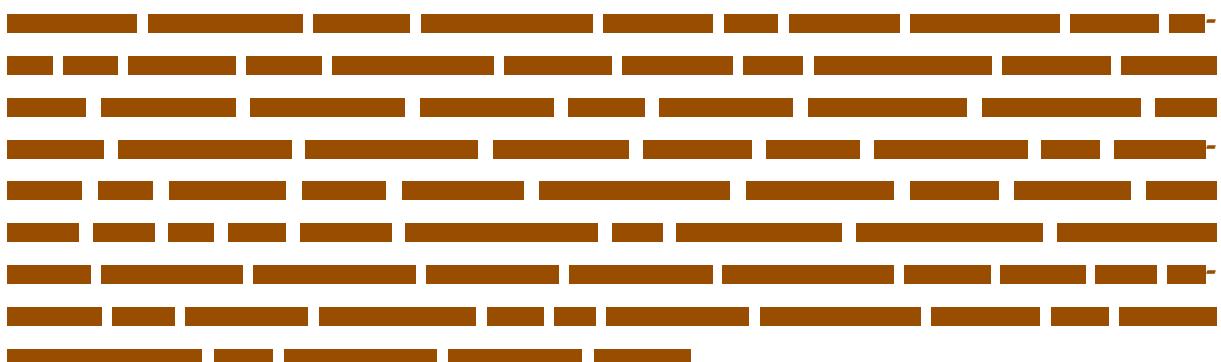
\starttexdefinition FakeOne #1#2#3#4#5
  \FakeScale{\FakeWidth / 400}
  \ruledhbox to 1tw {
    \llap {
      \ttxx
      \hpack {
        #3\enspace          % width
        \twodigits{#4}\enspace % kern
        #5\quad              % xscale
      }
    }
    \dorecurse {#2} {
      \Fake{#1}{#3}{#4}{#5}
      \unkern
      \space
    }
  }
\stoptexdefinition
```

```
\removeunwantedspaces
}
\stoptexdefinition
```

As a bonus we mention another ‘*trick*’ that sometimes gets applied: making spaces larger or smaller. In TEX this happens by applying stretch or shrink. We decided to stick to this model of spaces bound to fonts but not all fonts set up the spacing as we like. So, if one is desperate one can do this:

A series of 10 horizontal lines of fake text, each consisting of a sequence of short brown bars of varying widths. The lines are evenly spaced vertically, and each line represents a single line of text where the spacing between words has been manually adjusted by the user.

The previous fake text can also be typeset with `\spaceskipfactor 1500` but does it look better or worse? Our experience is that there is no reason to mess with the defaults, so this option is only there to prove a point.

A series of 10 horizontal lines of fake text, similar to the first but with a different font or style. Each line consists of a sequence of short brown bars of varying widths, showing the effect of the `\spaceskipfactor 1500` option on the typesetting.

