

# Interview of Frank Mittelbach

A combined interview of the L<sup>A</sup>T<sub>E</sub>X Project director

Frank Mittelbach

Gianluca Pignalberi

Dave Walden

**F**ree Software Magazine and the T<sub>E</sub>X Users Group (TUG (<http://www.tug.org/>)) both like to publish interviews. Recently, Gianluca Pignalberi of *Free Software Magazine* and Dave Walden of TUG both approached Frank Mittelbach about interviewing him. Rather than doing two separate interviews, Mittelbach, Pignalberi, and Walden decided on a combined interview in keeping with the mutual interests already shared by *Free Software Magazine* and TUG.

**DW: Frank, please start by telling us a bit about yourself and how you got involved with L<sup>A</sup>T<sub>E</sub>X.**

FM: I have lived with my family in Mainz (Germany) since the early eighties, i.e., by now the larger part of my life. Besides my primary hobby (typography), which can effectively be called my second job, I enjoy playing good board games, listening to jazz music, and reading (primarily English literature). Professionally I work for Electronic Data Systems where these days I'm responsible for concepts and implementation for remote monitoring and management of distributed systems and networks.

While I was studying Mathematics and Computer Science at the Gutenberg-University Mainz, I was first introduced to T<sub>E</sub>X and later L<sup>A</sup>T<sub>E</sub>X and, eventually, this got me interested in typesetting and in particular in research on algorithms for automated high quality typesetting.

During my student days in the eighties, a friend brought back a source tape from Stanford University containing something like T<sub>E</sub>X 1.1 and fascinating news about the qual-

ity of that program (back then we did our theses using a typewriter and either hand-pasting symbols in or, in case of some sophisticated IBM typewriter, changing the “ball” every couple of seconds). He tried to implement that program on the Multics system we had at the university and in fact succeeded—it probably was the first if not the only implementation of T<sub>E</sub>X on this operating system.

This way I got introduced to T<sub>E</sub>X and  $\mathcal{A}\mathcal{M}\mathcal{S}$ -T<sub>E</sub>X and typed my first paper, achieving beautiful results. The only catch was that back then the Stanford tape only contained

Figure 1: Frank Mittelbach



---

Almost Computer Modern fonts in 200 dpi resolution (and no METAFONT) and the only graphical printing device available to us had a resolution of 72 dpi. So the output we got was of the size of the following formula (or bigger):

$$\sum_{i=1}^n x_i$$

wonderful to plaster the walls with, but not necessarily suitable for handing in your thesis. As a result, my friend finally had to type his diploma thesis in the traditional way, despite his efforts.

Sometime afterwards I was asked by the department to install a commercial  $\text{T}_{\text{E}}\text{X}$  installation on our shiny new PCs and to give a series of lectures to students and professors on how to use it. And that distribution came with  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  2.08 and a loose-bound copy of the manual (which later became Leslie Lamport's book on  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ ).  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  compared to plain  $\text{T}_{\text{E}}\text{X}$  looked very good to me, but alas, when trying to produce any document with it, it died while loading the "article document style" due a lack of memory on those PCs. So my introduction to  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  stopped after I read the manual, and I was forced to develop my own  $\text{T}_{\text{E}}\text{X}$  macro package that implemented similar concepts while requiring less memory. A year later it became possible to actually use  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  at the department, and we could retire my macro package.

But this initial exercise gave me a good insight into the inner workings and concepts of a system like  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  and enabled me later to constructively criticize certain aspects of  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ —something that eventually led to Leslie passing on the development and maintenance of  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  to me.

**GP: Many of our readers are familiar with  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ , but for those who aren't, can you introduce  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  to them?**

FM:  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  is a batch-oriented typesetting system that uses the typesetting engine  $\text{T}_{\text{E}}\text{X}$  or one of its variants ( $\epsilon\text{-T}_{\text{E}}\text{X}$ ,  $\text{pdfT}_{\text{E}}\text{X}$ ,  $\text{Omega}$ ).

The  $\text{T}_{\text{E}}\text{X}$  program itself (developed by Professor Donald Knuth in the early eighties) is a programmable low-level typesetting engine, whose concepts and algorithms provide micro-typography<sup>1</sup> knowledge of highest quality in these

---

<sup>1</sup> Micro-typography is concerned with the detailed aspects of type and spacing, e.g., the kerning (shortening or enlarging space) between letters, generation and placement of ligatures, line breaking, etc. In contrast, macro-typography is concerned with larger structures, such as the design of headings, lists, or pages.

days when this knowledge is slowly declining due to the fact that more and more authors are forced to become their own designer and typesetter without proper training.  $\text{T}_{\text{E}}\text{X}$  is especially known for its excellent paragraph breaking algorithm and for its math formula typesetting capabilities, both of which are unsurpassed even though the program and its algorithms have been freely available for more than twenty years.

$\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  is a macro package written for the  $\text{T}_{\text{E}}\text{X}$  engine which allows the user to step back from the low-level formatting capabilities of  $\text{T}_{\text{E}}\text{X}$  by providing higher-level interfaces that give the author the ability to mark up the text with logical markup rather than procedural markup (e.g., specifying that something is a list or a section, rather than stating that something should be set in a bold typeface with a little space above and below). The actual transformation of a  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  source into a typeset document is done with the help of "style sheets" and configuration adjustments that allow even radical changes to the design and layout in a consistent manner without touching or changing the source<sup>2</sup>.

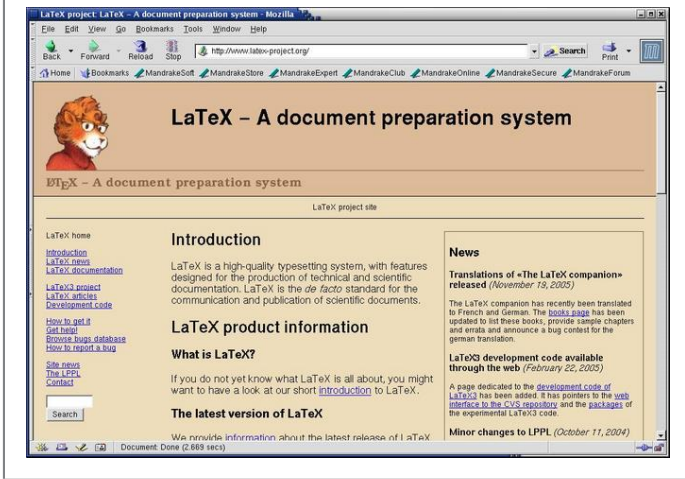
Historically speaking,  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  was largely influenced by a system called Scribe (by Brian Reid). In turn,  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ 's concept of logical markup was quite influential on HTML and various SGML/XML DTDs, as were its approaches for turning such logical markup into visual representation.

One of the differences between  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  and many other similar approaches is that the  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  language is in fact a community development: new packages that augment (or modify)  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ 's markup and typesetting functionalities are constantly appearing, so that these days  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  offers typesetting solutions for nearly every subject domain—as diverse as game typesetting (such as chess, go, or crossword puzzles), chemical formulas, or music. Another important difference is that, although  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  brought the concept of logical markup to a larger audience, it also provides ways to fine-tune the results (essentially providing interfaces to procedural markup), acknowledging the fact that no automated transformation of logical markup into a visual representation is able to automatically resolve all problems produced by the physical restrictions of the output format (e.g., line width or page size). While in certain applications such fine tuning adds no value (like database content publishing

---

<sup>2</sup> Well, ideally, but see below.

Figure 2: The L<sup>A</sup>T<sub>E</sub>X Project home page



where full automation is required), it is crucial for typesetting high quality books and journal articles.

**GP: How many people are officially part of the L<sup>A</sup>T<sub>E</sub>X Project? And how would you define the “L<sup>A</sup>T<sub>E</sub>X project”?**

FM: The L<sup>A</sup>T<sub>E</sub>X Project Team is a fairly small (slowly changing) group of people who look after the L<sup>A</sup>T<sub>E</sub>X kernel and a small number of core packages that provide a stable basis for a huge number of constantly evolving packages and add-ons. Providing and guarding a stable core is (although not necessarily popular with everyone) an important part in keeping L<sup>A</sup>T<sub>E</sub>X alive as a language for document exchange. Current and past members of the team include Javier Bezos, Johannes Braams, David Carlisle, Michael Downes, Denys Duchier, Robin Fairbairns, Morten Høgholm, Alan Jeffrey, Thomas Lotze, Chris Rowley, Rainer Schöpf, and Martin Schröder with varying degrees of involvement.

Historically, the project took over maintenance and development of L<sup>A</sup>T<sub>E</sub>X 2.09 from Leslie Lamport in 1991. At one time the system was split into several incompatible variants that often prohibited successful processing by L<sup>A</sup>T<sub>E</sub>X at one site of documents created by L<sup>A</sup>T<sub>E</sub>X from a different site, even though system independence was originally one of the important goals of L<sup>A</sup>T<sub>E</sub>X as a documentation language for the scientific community. Another goal for the team was to address apparent deficiencies in the concepts of L<sup>A</sup>T<sub>E</sub>X 2.09. The project team addressed both issues in the early nineties with L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub>, which provided a stable and consolidated

platform offering further development possibilities outside the kernel code.

Although L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub> already addressed most, if not all, of the deficiencies identified in the first decade of L<sup>A</sup>T<sub>E</sub>X 2.09 use, it was originally thought that L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub> would only be an intermediate step towards a L<sup>A</sup>T<sub>E</sub>X 3 version. But over time it became clearer and clearer that the remaining open questions could not be adequately resolved within the constraints of: a) T<sub>E</sub>X as the underlying formatting engine, and b) no changes in the fundamental concepts deployed in L<sup>A</sup>T<sub>E</sub>X.

As a result, most of the efforts in the recent years by members of the L<sup>A</sup>T<sub>E</sub>X team have gone into research on features desirable for the underlying formatter engines and into development of experimental languages and concepts for a designer’s interface to typesetting—a level of abstraction that is largely missing from today’s L<sup>A</sup>T<sub>E</sub>X (which currently often requires T<sub>E</sub>X programming).

So these days, one definition of the L<sup>A</sup>T<sub>E</sub>X project would be that it works on providing the foundation for the core concepts and implementation of a new typesetting system that is based upon the good aspects of L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub> (e.g., logical markup, extensibility), but, on the code level, it’s not necessarily focused on providing compatibility. At the document syntax level, the situation is clearly different, as reuse of older documents is certainly an important goal. But, even there, the main focus will be on clean concepts; and, as a result, compatibility may be restricted in certain cases to providing support for automated conversion.

**DW: You have a long history in the world of T<sub>E</sub>X for collaborative work—famously, with Rainer Schöpf, in the early days of L<sup>A</sup>T<sub>E</sub>X. You must enjoy working collaboratively. Tell us a bit about your approaches to collaborative work.**

FM: I do indeed like to collaborate; and, over the years, worked successfully with many different people (on various topics and in different subject domains). For me, the main value of collaboration is during the development of ideas, which, in my experience, are best produced in an open exchange. My mental picture here is a table tennis or similar game which only develops if one directly reacts to whatever your counterpart thinks of and “picks up the ball as played”. People who have worked with me know that I like white board drawing sessions (I do need to visualize while I play along) and brainstorming and mind mapping methods.



---

But I'm also a stickler for details and can spend a lot of energy and effort in actually finishing something (to my own satisfaction), when I consider it worthwhile. Collaboration on that level—after the initial concept and design development work has finished and the nitty gritty detail work starts—normally takes one of two forms: either I restrict myself largely to mentoring and let others work on actual implementations, or I put so much energy into a certain task that it outweighs other people's involvement by a large factor. My base motto here is “Es gibt nichts Gutes, außer man tut es” (free translation: Nothing good will come into existence unless you actually do it) by Erich Kästner, which, at least in the German language, nicely rhymes.

A lot of collaboration necessarily happens via email (due to living in different countries, etc.), but I find it extremely valuable to interrupt this method of working at irregular intervals, with face-to-face meetings to flesh out ideas and make them concrete enough to go ahead for a while in semi-isolation, with only email and/or phone calls as the means of “direct” communication. This also explains why most of the more fundamental work, associated with Rainer Schöpf and me, dates from the time when we both studied at university and had a chance for a more regular exchange of ideas in front of white boards (drinking gallons of tea).

In general, I think that Frederick Brooks is right when he argues in *The Mythical Man-Month* [4] that to run a successful software project, you need a fairly small and structured team that is responsible for making the final design decisions. Large scale “committee” design only leads to bad results, by compromising too much between different factions, or by incorporating incompatible design concepts.

**GP: T<sub>E</sub>X is widely considered the best typesetting system, but professional typesetters seem to prefer using commercial, visual software. Why would you advise them to use L<sup>A</sup>T<sub>E</sub>X instead of another system? Or, conversely, why not?**

FM: There is no doubt that T<sub>E</sub>X has superior qualities in a number of areas when compared to other typesetting systems (e.g., paragraph breaking quality, or math formula presentation, etc.). But, it was designed as a batch processing program; that is, it does not allow for direct interaction with the user. In WYSIWYG typesetting systems the user can make visual corrections which are then instantaneously reflected, while with T<sub>E</sub>X (or L<sup>A</sup>T<sub>E</sub>X) you have to modify your

source, and then reprocess and check that your correction produced the desired result.

Especially in the last stage of book production, T<sub>E</sub>X's tendency for making far reaching changes to achieve high quality can actually be a hindrance rather than an asset (at least, if you don't account for it and adjust your working method). For example, due to global optimization in paragraph breaking, a *removal* of a single word in one paragraph will usually result in a complete reflow of the whole paragraph. This might, in fact, make the paragraph one line *longer*, if T<sub>E</sub>X decides that a slightly looser setting of all lines produces the best possible solution. If this happens when correcting last-minute typos, in an otherwise finally formatted document, it can be rather annoying.

Most graphic designers and professional typesetters are used to working visually with immediate feedback and control. So, for them, systems like L<sup>A</sup>T<sub>E</sub>X appear difficult to handle and they don't see any benefit in this unfamiliar working model.

---

L<sup>A</sup>T<sub>E</sub>X is a batch-oriented typesetting system that uses the typesetting engine T<sub>E</sub>X or one of its variants (ε-T<sub>E</sub>X, pdfT<sub>E</sub>X, Omega)

---

Whether or not the use of L<sup>A</sup>T<sub>E</sub>X would be advisable really depends on the job at hand, and cannot be answered without context. In a nutshell, I would suggest using L<sup>A</sup>T<sub>E</sub>X or a similar system whenever one or more of the following factors play an important role in the job:

- user's preference is to think in logical structures
- designs that require consistency
- documents whose designs are not yet fully defined or that need to be presented in several layouts in parallel
- documents that require high-quality paragraph breaking
- documents that contain heavy mathematics
- automatically produced content (e.g., from databases)
- long material

On the other hand, the following factors move the balance towards using a (good quality!) visually oriented system:

- user's preference is to think in visual structures

- user not at all comfortable working with programming languages (a high-level front end for T<sub>E</sub>X, e.g., L<sup>A</sup>T<sub>E</sub>X or ConT<sub>E</sub>Xt, helps here but...)
- designs that require a lot of visual flexibility rather than consistency (e.g., headings are designed one-off according to nearby objects)
- designs that require text to flow around arbitrary shapes (T<sub>E</sub>X is simply not designed for this)
- designs that change the horizontal measure from column to column
- short material

What actually tips the balance may differ in different circumstances. For example, in *The L<sup>A</sup>T<sub>E</sub>X Companion* [9], with its nearly 1000 in-line examples, it was an enormous plus to be able to redesign example layout without touching the individual examples.

**DW: Is there any relationship between TUG, DANTE, or any of the other T<sub>E</sub>X users groups and the L<sup>A</sup>T<sub>E</sub>X Project, and how might the user groups help maintain the viability of T<sub>E</sub>X and its derivatives?**

FM: The relationship between the L<sup>A</sup>T<sub>E</sub>X project and T<sub>E</sub>X user groups can probably be best described as loose and informal. Several project team activities have been supported, in one way or the other, by a user group (e.g., by providing meeting space at a conference). However, most L<sup>A</sup>T<sub>E</sub>X team activities have drawn on non-user-group resources such as support from ZDV (the computing laboratory at the Gutenberg University Mainz), royalty payments from *The L<sup>A</sup>T<sub>E</sub>X Companion*, and to a small extent from individual user contributions. The biggest joint venture with a user group was probably the development of a model for a set of extended math fonts, where the French user group financed a student for three months to work with me on this topic. The outcome of this work [13] is now finally bearing some fruit as it helped in developing the STIX fonts.

However, the user groups are extremely important to projects such as the L<sup>A</sup>T<sub>E</sub>X Project in that they provide a research framework for contacts and face-to-face discussion at conferences and journal publications. This aspect of providing a research framework cannot be underestimated. And, if the user groups would become unable to provide it, it might eventually result in the death of the community. I certainly enjoy (and I'm sure so do others) the fruitful exchange that is only possible in such a framework.

The question about what the user groups can do long term to sustain the viability of T<sub>E</sub>X and its derivatives, is difficult to answer—I'm unfortunately not sure the user groups themselves will survive in the long run. The role of the user organizations has changed over the last two decades. In the beginning, just getting a T<sub>E</sub>X system installed was a major effort, and user groups were formed by interested people to help each other, exchange knowledge and ideas, and to support development effort. Back then, the role of the user groups was fairly clear and the benefit for each member was immediately visible (e.g., obtaining information otherwise not available, getting help, etc.).

Over time, access to a T<sub>E</sub>X installation and all its accompanying goodies (such as L<sup>A</sup>T<sub>E</sub>X packages, etc.) has drastically changed—nowadays installations are prepackaged, access to all software is available in large archives, and there is much more documentation available. As a result, the typical T<sub>E</sub>X/L<sup>A</sup>T<sub>E</sub>X user has no need to understand the underlying mechanisms and (unfortunately) isn't any longer interested in sharing in their development. The users have largely changed from actively participating members of a group of like-minded people, excited by the possibility of doing high-quality typesetting, to consumers of a "finished product" who get very upset if the product doesn't do precisely what they want it to do. For this new kind of user, the users groups don't play an important role since, at this point in time, the users groups have no resources to actually help individuals with their problems.

This is somewhat ironic, since it was largely members of the user groups that initiated all the changes that now appear to be leading to the downfall of the user groups' accepted mission and reason for existence in the eyes of the average user.

In theory, I think the best way that users groups could help these days would be in the following areas:

- recruiting and providing the resources that keep the "product" alive and well-maintained
- providing a suitable forum for the active development community
- obtaining and managing research funds
- attracting new users to broaden the base

Unfortunately, none of this is easily achievable. It would need a noticeable amount of capital (and resources), beyond

what is currently available to the groups, and it isn't clear that this—as a charter—would attract enough new members, who could then share those costs. After all, to most people the “product” and its support appear to be available free of charge. So, today's consumer thinking is: “Why pay a (substantial) recurring membership fee, when all that is needed is to connect to the internet and ask a question on `comp.text.tex` or download some software from CTAN? I don't go to those conferences ‘they’ go to, so why should I finance ‘them’?” It would be necessary to break that thinking and make people understand why the user groups are, nevertheless, beneficial for them. But, unfortunately, many people take a free lunch if they can get it, without considering the consequences.


**GP: Your project generated a license: the L<sup>A</sup>T<sub>E</sub>X Project Public License. Thanks to the last modifications to it, L<sup>A</sup>T<sub>E</sub>X can really be considered free software. How did this fact improve the diffusion of such a tool?**

FM: To be honest, I always considered and still consider L<sup>A</sup>T<sub>E</sub>X *real free software* regardless of the license under which it was distributed in the past or is distributed now. The term “free” definitely has different meanings for different people, and I don't necessarily agree with the understanding of some people that their freedom to be able to arbitrarily change things without any restrictions, should be considered a more important good than the right of others to get what they expect, when they use a certain product.

L<sup>A</sup>T<sub>E</sub>X is not just a single user product, but a language being used for communication of information, and, one of the important points here is that it enables processing a document at different sites with identical results, provided that the same version of L<sup>A</sup>T<sub>E</sub>X is used. This is a feature a large proportion of the community is relying on, so the original LPPL (L<sup>A</sup>T<sub>E</sub>X Project Public License) [3] in a nutshell said: do whatever you like with file X but if you change it (i.e., modify its behavior in the system) change its name to something else, so that people relying on the communication feature of L<sup>A</sup>T<sub>E</sub>X will not be affected by your modification.

Technically, this allowed for any modification and any desired change, but it gave the people using L<sup>A</sup>T<sub>E</sub>X a conscious choice to apply a changed version to their documents or not. In some cases, it would have meant some extra effort for the person doing modifications, but on the whole I feel it provided a nice balance between the people who think “free”

Figure 3: The L<sup>A</sup>T<sub>E</sub>X Companions: the definitive boxed set



means their right to change what they like, and people who think they have a right to a reliable means of communicating information.

However, some developers in the free software community think that such a simple rule restricts their rights too much (not being able to change things in arbitrary ways, including ways that hide the modification to later users—even if that is not the intention). So, a discussion started about whether or not such a rule makes software non-free—the main obstacle for many being the requirement to change names if you change content. Clearly, this requirement is quite different from those posed by the majority of free software licenses, but then those licenses have been written with quite a different software model in mind (one where the focus lies on individual software components where differences at different sites do not restrict the usability of the software)<sup>3</sup>. But, since we were not interested in enforcing a name change per se (even though we still think that it mediates nicely between all different needs), I entered a longish discussion with `debian-legal` and, as a result, we came up

<sup>3</sup> For a discussion of why we think that something like GPL is not a good licensing model for free languages, which is one aspect of L<sup>A</sup>T<sub>E</sub>X, see [2].

with a new license which softened this requirement while still preserving the community need for stability and reliability. In essence, I think the new license is better in many parts (and I'm very grateful to some folks from debian-legal for helping achieve this), but it's also, perhaps unnecessarily, more complex than it could have been in other parts. In the  $\TeX$  world, the original license was trying to codify what was standard and accepted behavior, i.e., when you changed or improved a package you called the result something else so that older documents would compile as expected, while newer ones could make use of extended or changed features and both could co-exist.

So, did the license improve the diffusion of the tools? As far as the  $\TeX$  world as such is concerned, I would say "no", as even the original license was already simply codifying what most people thought to be a good model for software in the  $\TeX$  domain. As to the wider world of free software in general, the modifications probably helped people to understand that  $\LaTeX$  and friends are also "free" software, and provided common ground for some understanding that different usage requirements may need somewhat different interpretations of "free" to be useful.

**GP: Are you, were you, or will you be involved in other free software projects?**

FM: The answer to this probably depends on the definition of the terms "project" and "involvement". Many of my interests these days are of a more theoretical nature and will not necessarily directly lead to software, or not to software where I will directly participate in implementations. Those projects where I'll most likely participate could be labeled under the broad heading of  $\TeX/\LaTeX$  development. For example, just a couple of months ago Hàn Thế Thành (the developer of pdf $\TeX$ ), Morten Høgholm and I spent a productive weekend at my home working on ideas for grid typesetting<sup>4</sup>. So, even though I most likely won't participate in actual implementations, there is and will be involvement in projects outside of  $\LaTeX$ . And who knows, as I'm doing completely different work in my professional life, perhaps that too will one day lead to one free software package or another in that area.

**DW: You, among others, have written about the need to move beyond the limitations of  $\TeX$  and suggested**

**improved approaches. In your biography in *The  $\LaTeX$  Companion* [9], you say that you want to work at bringing extensions such as Omega and  $\varepsilon\text{-}\TeX$  together as a base for an actual  $\LaTeX$  3. Obviously, you have a track record for accomplishing big, complex,  $\TeX$ y projects. Do you envision getting involved with something like the ex $\TeX$  project or starting your own low-level implementation project for an improvement on  $\TeX$ ?**

FM: When I wrote *E- $\TeX$ : Guidelines for future  $\TeX$  extensions* [7] in 1990, the time was not yet ripe for improving  $\TeX$ , and many people actually considered it an affront to Don that I suggested there could be something worth improving in his product. (I remember, for example, public musing about strange theories from unknown and obscure German typographers. Well, those "strange theories" had been suggested to me by none other than Hermann Zapf, who, though German, may not be precisely called unknown let alone obscure.) But, be that as it may, what I was challenging in that paper was the typesetting quality at the micro-typography level. However, as  $\TeX$  was technically so much better than anything else at that time, my challenge was probably premature, and it took nearly a decade before the first real experiments were conducted on that level and things moved ahead in that domain (largely with the development of experimental versions of pdf $\TeX$  but also experimental code by others, e.g., Matthias Clasen).

When discussing improvements to  $\TeX$  one needs to distinguish three largely disjunct areas. The first is the area of the programming language and the fact that this language is incomplete and for certain tasks difficult to use (or, as some people state, "a mess"). On that level (without diverting from the fundamental paradigms of  $\TeX$ ), extensions like  $\varepsilon\text{-}\TeX$ , and to some extent Omega, tried to ease the programming task by providing missing primitives that bridge the obvious gaps in the base language. But, since such additional functionality was only easing the programmer's life without actually improving the typeset results and the function of (nearly) all the new primitives could be achieved with some extra effort in the base language, we decided to stay away from them in the  $\LaTeX$  development, as their use would have resulted in a  $\LaTeX$  version that would then only run on a small fraction of the installations without any practical gain for the user. The  $\LaTeX$  Project together with people from Con $\TeX$ t and  $\varepsilon\text{-}\TeX$  actually made some effort to produce an enriched syntax definition for  $\varepsilon\text{-}\TeX$  [5] that we

<sup>4</sup> Grid typesetting describes designs that are based on an invisible underlying grid restricting the placement (and size) of objects.



---

thought would provide enough benefits to switch to a  $\TeX$  successor implementing this extended set. Sadly, shortly after this proposal, work on  $\varepsilon\text{-}\TeX$  effectively came to a standstill, and none of this was ever implemented. Nevertheless, something has changed since then: the installed base of  $\varepsilon\text{-}\TeX$ -enabled installations did grow beyond critical mass (largely because of  $\text{pdf}\TeX$  which included the  $\varepsilon\text{-}\TeX$  extensions), so that some time ago, the  $\LaTeX$  Project officially announced that it will base future  $\LaTeX$  versions on this extended set of primitives—and recently started to actually produce code that made use of these extended features (although so far only outside the kernel code). In essence, we never wanted to go away from being able to have  $\LaTeX$  run “out of the box” on a large base of installed interpreter programs, and valued this higher than a potentially easier or better adjusted programming language that nearly nobody could use<sup>5</sup>. So, instead of only trying to influence the  $\TeX$  language by extending it, I and some others in the  $\LaTeX$  project also worked from the inside by developing the “Experimental  $\LaTeX$  programming language” [6]. This was done over several prototypes (the first already done in 1993 or so). The current version is something we think can be used successfully and we have started to provide the first public packages in this language [1].

The second area is the one dealing with micro-typography issues, e.g., those that I was mainly concerned with when discussing shortcomings of  $\TeX$  in [7]. In this area my involvement was largely confined to initiating work by others. The third area is the one that concerns itself with the generally open and unsolved questions of computer typography, e.g., models for representation of logical [8] and visual content material [12]; transformation between logical and visual representation using automated methods that nevertheless provide highest quality according to a defined metric [10] to give some examples. Part of that research is to understand and codify typography rules and to develop concepts and algorithms that can be driven by parameterized rules, e.g., to produce high-quality float placement.

---

<sup>5</sup> I learned that this is a critical factor when we tried to introduce  $\LaTeX 2_{\varepsilon}$  in 1994 which required the installation of T1 encoded fonts (i.e., fonts containing characters with diacritics). The switch to the new system nearly collapsed, because users in the US saw absolutely no benefit in a system that contained all these useless characters only needed by Europeans. Fortunately enough,  $\LaTeX 2_{\varepsilon}$  had other benefits that eventually won over nearly all  $\LaTeX$  users, but it was a close shave.

Do I envision starting my own low-level implementation project to improve on  $\TeX$ ? Most certainly not, but I do envision getting (re)involved with the developments currently happening and hope to bring some of those developments together. Whether this will be in a project like  $\text{ex}\TeX$  or  $\text{pdf}\TeX$  is largely irrelevant. At this point in time there are still many unresolved questions, and it is still the time for experiments (which may happen in different projects in parallel), but one important goal would be to bring the various developers together to talk to each other about their ideas and the concepts behind the ideas.

---

System independence was originally one of the important goals of  $\LaTeX$  as a documentation language for the scientific community

---

**GP: You mentioned two operating systems  $\LaTeX$  was ported to, and we know it runs on several free and non-free (whether commercial or not) OSs. What kind of OSs and programs do you mainly use? And why?**

FM: To be precise,  $\LaTeX$  is interpreted so it is not software that needs porting to any OS;  $\LaTeX$  runs everywhere where  $\TeX$  has been ported to—and  $\TeX$  to my knowledge has been ported to more or less every operating system ever in existence (with the exception of something like the Palm OS), i.e., I have used it on Mainframes, VMS, Unix, Multics, and Windows.

I use both free and commercial operating systems; it largely depends on the environment and the task at hand. At home, I run mainly Linux with VMware to access certain programs only available on Windows. On my laptop the situation is reversed: I run XP native and use cygwin for a decent command line environment with all the benefits of a good Unix system. My favorite editor is Emacs, which I use on nearly every platform. I like to structure things using mind maps and for this the only really good program I have found is commercial and works only on Windows—it is one of the reasons that these days I use Windows fairly regularly<sup>6</sup>.

In the professional world, where I earn my living, the predominant OS on the desktop is Windows; and in the server

---

<sup>6</sup> I have recently reevaluated the free program freemind and it has greatly improved, so this may change again.



---

world you'll find commercial Unix variants but also a growing number of Linux servers.

-----

**L<sup>A</sup>T<sub>E</sub>X runs everywhere where T<sub>E</sub>X has been ported to—and T<sub>E</sub>X to my knowledge has been ported to more or less every operating system ever in existence (with the exception of something like the Palm OS)**

-----

What I use depends largely on the task at hand: for some, e.g., project management, the Windows world simply offers the better tools; in others, free software (running on commercial or free OSs) provides better quality, or features otherwise not available. Examples would be Perl, Apache, CVS, subversion, and others.

**DW: I remember Knuth saying that writing *The T<sub>E</sub>Xbook* led to hundreds of changes in T<sub>E</sub>X, because he was forced to explain things to the reader, and when he couldn't, he changed the program. You have written a number of books and have been a major developer. Has your work as an author influenced your work as a developer, or vice versa?**

FM: I think Don is absolutely right in making such a statement: I think it is extremely valuable to combine the development of software (actually anything) with the task of writing about it in some way. Trying to explain to others the functions and concepts behind a creation helps a lot in finding out whether or not something is going to work in practice. If you can't explain it or if the explanation turns out to be horribly complicated, then there is something fundamentally wrong with your creation and you should return to the drawing board.

Very important here is that one doesn't stop at simply documenting functions or menu items (though that is a start), but effectively tries to document the usage flow and the reasons why one would do things in one way or the other. Often enough (with free software as well as commercial software), you find only rudimentary documentation that tells you that such and such feature exists, but never explains why one would want to use the feature in the first place. That type of documentation, while necessary, will not help in improving your tool (and, often enough, it turns out that such features

only got implemented because they were easy to add without providing any real benefit).

So yes, documenting ideas and work flows has always been an integral part for me of developing and/or improving my own software as well as software from others. In *The L<sup>A</sup>T<sub>E</sub>X Companion* [9], for example, a good proportion of what I describe is software developed by others, and the process of trying to explain how to use this software and finding good usage examples led in many cases to improvements in syntax or features after some discussions with the authors.

Therefore, my advice to developers is to always try their hands at documenting their own creations, or at least find somebody who does it for them (starting from the initial development!)—and carefully evaluate the findings from this process: it will result in noticeable improvements in the product.

**DW: Thank you very much, Frank, for taking the time to do this interview with us. Your insights about T<sub>E</sub>X, L<sup>A</sup>T<sub>E</sub>X, and the development and diffusion of complicated systems in a distributed development environment are fascinating. I believe the readers of TUG's Interview Corner will agree.**

**And thank you, Gianluca, for agreeing to let me share this interview with you. It has been a pleasure to work with you.**

**GP: Thank you, Frank, for giving *Free Software Magazine* readers a very well explained essay about an important piece of free software. Moreover, your LPPL explanation clarified some obscure points in a previous article [11]. And thank you, Dave, for giving me the possibility of doing this combined interview, which is much more interesting than a "normal" interview.**

FM: Thanks, Dave and Gianluca, for conducting this interview in the way it was done. I enjoyed seeing it unfold, question after question—despite the time it took (my fault) and the fact that we live far apart, it felt like doing a live interview face to face, which, I think, is the way it should be.

## Bibliography

[1] L<sup>A</sup>T<sub>E</sub>X 3 Project Team. The L<sup>A</sup>T<sub>E</sub>X Project CVS Repository. Available at the L<sup>A</sup>T<sub>E</sub>X Project website (<http://www.latex-project.org/cgi-bin/cvsweb/>).

[2] L<sup>A</sup>T<sub>E</sub>X 3 Project Team. Modifying L<sup>A</sup>T<sub>E</sub>X. Available with all L<sup>A</sup>T<sub>E</sub>X distributions, December 1995.

[3] L<sup>A</sup>T<sub>E</sub>X 3 Project Team. The L<sup>A</sup>T<sub>E</sub>X Project Public License (Version 1.3). Available at the L<sup>A</sup>T<sub>E</sub>X Project website (<http://www.latex-project.org/lppl/>), December 2003.

[4] Frederick P. Brooks, Jr. *The Mythical Man-Month; Essays on Software Engineering*. Addison-Wesley, Boston, Massachusetts, 2nd edition, 1995. ISBN 0-201-83595-9.

[5] David Carlisle. Notes on the Oldenburg  $\epsilon$ -T<sub>E</sub>X/L<sup>A</sup>T<sub>E</sub>X 3/ConT<sub>E</sub>Xt meeting. Available at the L<sup>A</sup>T<sub>E</sub>X Project website (<http://www.latex-project.org/papers/etex-meeting-notes.pdf>), 1998.

[6] David Carlisle, Chris Rowley, and Frank Mittelbach. The L<sup>A</sup>T<sub>E</sub>X 3 Programming Language—a proposed system for T<sub>E</sub>X macro programming. *TUGboat*, 18(4):303–308, December 1997. ISSN 0896-3207. Available at the T<sub>E</sub>X User Group website (<http://www.tug.org/TUGboat/Articles/tb18-4/tb57rowl.pdf>).

[7] Frank Mittelbach. E-T<sub>E</sub>X: Guidelines for Future T<sub>E</sub>X Extensions. *TUGboat*, 11(3):337–345, September 1990. ISSN 0896-3207. Available at the T<sub>E</sub>X User Group website (<http://www.tug.org/TUGboat/Articles/tb11-3/tb29mitt.pdf>).

[8] Frank Mittelbach. Language Information in Structured Documents: Markup and rendering—Concepts and problems. In *International Symposium on Multilingual Information Processing*, pages 93–104, Tsukuba, Japan, March 1997. Invited paper. Republished in *TUGboat* 18(3):199–205, 1997. Available at the T<sub>E</sub>X User Group website (<http://www.tug.org/TUGboat/Articles/tb18-3/tb56lang.pdf>).

[9] Frank Mittelbach and Michel Goossens. *The L<sup>A</sup>T<sub>E</sub>X Companion*. Tools and Techniques for Computer Typesetting. Addison-Wesley, Boston, Massachusetts, 2nd edition, 2004. ISBN 0-201-36299-6. With Johannes Braams, David Carlisle, and Chris Rowley.

[10] Frank Mittelbach and Chris Rowley. The pursuit of quality: How can automated typesetting achieve the highest standards of craft typography? In C. Vanoirbeek and

G. Coray, editors, *EP92—Proceedings of Electronic Publishing, '92, International Conference on Electronic Publishing, Document Manipulation, and Typography*, Swiss Federal Institute of Technology, Lausanne, Switzerland, April 7–10, 1992, pages 261–273, New York, 1992. Cambridge University Press. ISBN 0-521-43277-4.

[11] Gianluca Pignalberi. The L<sup>A</sup>T<sub>E</sub>X Project Public License. *Free Software Magazine*, (7):52–54, 2005. ISSN 1746-8752. Available at the Free Software Magazine website ([http://www.freesoftwaremagazine.com/free\\_issues/issue\\_07/tex\\_license/](http://www.freesoftwaremagazine.com/free_issues/issue_07/tex_license/)).

[12] Chris A. Rowley and Frank Mittelbach. Application-independent representation of multilingual text. In Unicode Consortium, editor, *Europe, Software + the Internet: Going Global with Unicode: Tenth International Unicode Conference, March 10–12, 1997, Mainz, Germany*, San Jose, CA, 1997. The Unicode Consortium. Available at the L<sup>A</sup>T<sub>E</sub>X Project website (<http://www.latex-project.org/papers/unicode5.pdf>).

[13] Justin Ziegler. Technical report on math font encoding (version 2). Technical report, L<sup>A</sup>T<sub>E</sub>X 3 project, June 1994. Available from CTAN: the CTAN website (<http://www.ctan.org/info/ltx3pub/l3d007.tex>).

## Copyright information

© 2006 Frank Mittelbach, Gianluca Pignalberi, Dave Walden

Verbatim copying and distribution of this entire article is permitted in any medium without royalty provided this notice is preserved.

## About the authors

**Gianluca Pignalberi** writes articles for *Free Software Magazine* and developed the L<sup>A</sup>T<sub>E</sub>X class with which the magazine is published.

**Dave Walden** developed the web publication software for TUG's on-line publication, *The PracT<sub>E</sub>X Journal*, and coordinates TUG's Interview Corner.

This is a **high resolution, color** version of the article you downloaded.  
You are kindly invited to subscribe!

Subscriptions are free, and every subscriber receives our fantastic weekly newsletters — which are in fact fully edited articles about free software.

Please click here to subscribe:

<http://www.freesoftwaremagazine.com/subscribe>