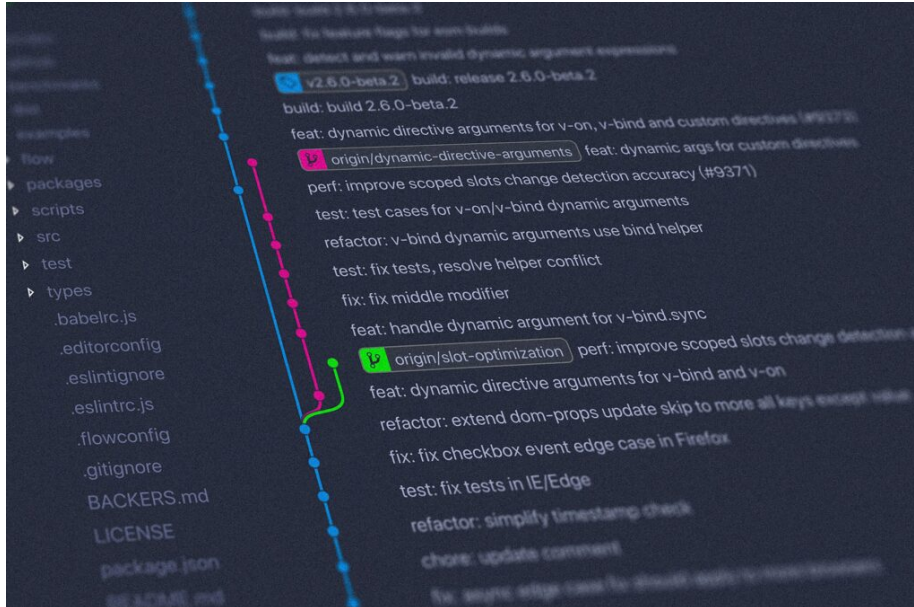


Issue 066: Version Control

Adrian Kosmaczewski

March 4th, 2024



Welcome to the sixty-sixth issue of *De Programmatica Ipsum*, about *Version Control*.

In this edition:

- We analyze the rise in popularity of Git¹ and how it eclipsed everything else.
- In the Library section², we review “Pragmatic Version Control Using Git” by Travis Swicegood³.
- In our Vidéothèque section⁴, we watch Linus Torvalds explaining Git⁵ in a 2007 Google TechTalk.

We would like to thank our patrons who generously contribute every month (or have contributed in the past) to our work and help us run this magazine. Thank you so much! In alphabetical order: Adam Guest, Adrian Tineo Cabello, Benjamin Sheldon, Christopher Nascone, Colin Powell, Franz Lucien Moersdorf, Guillermo Ramos Álvarez, Jean-Paul de Vooght, Dr. Juande Santander-Vela, Patryk Matuszewski, Paul Hudson, Quico Moya, Roger Turner, and Szymon Licau.

Enjoy this issue! Please subscribe to our free newsletter⁶ to stay updated about new releases, share the articles on social media, or contribute⁷ if you would like to support our work with a donation via Liberapay⁸.

¹<https://deprogrammaticaipsum.com/twenty-years-is-nothing/>

²<https://deprogrammaticaipsum.com/category/library/>

³<https://deprogrammaticaipsum.com/travis-swicegood/>

⁴<https://deprogrammaticaipsum.com/category/videotheque/>

⁵<https://deprogrammaticaipsum.com/linus-torvalds/>

⁶<https://deprogrammaticaipsum.com/newsletter/>

⁷<https://deprogrammaticaipsum.com/contribute/>

⁸<https://liberapay.com/>

Cover photo by Yancy Min⁹ on Unsplash¹⁰.

⁹https://unsplash.com/@yancymin?utm_content=creditCopyText&utm_medium=referral&utm_source=unsplash

¹⁰https://unsplash.com/photos/a-close-up-of-a-text-description-on-a-computer-screen-842ofHC6MaI?utm_content=creditCopyText&utm_medium=referral&utm_source=unsplash

Twenty Years Is Nothing

Adrian Kosmaczewski

March 4th, 2024



In a previous edition of this magazine¹, we argued that English was so pervasive in our industry, nobody even questioned its use anymore. The same can be said of Git². It is difficult to imagine that merely twenty years ago, the landscape of source control tools was more diverse, and the choice of one such tool was much more complicated than today. Actually, Git was not even on the map yet. Before debating whether the hegemony of Git is good or bad, let us go back in time for a little while.

In one of the most famous tangos³ of all time, Carlos Gardel⁴ famously sings

To feel... that life is a breath of fresh air,
that twenty years is nothing,
that, feverish, the gaze
wandering in the shadows
seeks and names you.

Twenty Years Ago

The second edition of “Code Complete” by Steve McConnell⁵ was published in 2004. On page 668 of this massive 900-page volume, we find the *only* reference to the subject of source control in the whole book: about three quarters of a page long. Nothing else. ChatGPT

¹<https://deprogrammaticaipsum.com/the-winner-takes-it-all/>

²<https://git-scm.com/>

³[https://en.wikipedia.org/wiki/Volver_\(song\)](https://en.wikipedia.org/wiki/Volver_(song))

⁴https://en.wikipedia.org/wiki/Carlos_Gardel

⁵<https://deprogrammaticaipsum.com/steve-mcconnell/>

can easily summarize all of it with one phrase: “Version control software is good and brings several big benefits.” Not a lot to phone home about. We are really far from GitOps at this point.

That same year, almost exactly 20 years ago at the time of the publication of this article, Subversion 1.0⁶ saw the light of day. What was Subversion? Probably the shortest lived idea-with-good-intentions in the history of computing. See, Subversion (or svn) was supposed to be a better CVS (no, not the pharmacy⁷, but this other thing⁸). “Better than CVS” meant, back in those days, to be transactional (databases⁹, anyone?) and to have a somewhat better support for branches. We did not have higher ambitions back then, kids.

Linus Torvalds, however, did have higher ambitions. In 2004, the Linux kernel developers got in an increasingly strong disagreement over the use of BitKeeper¹⁰, the proprietary, distributed version control system used to manage the kernel source code. So, what is a developer to do? Well, Linus has a tradition of writing the software everybody needs and nobody wants to start. He also has a tradition of naming things after himself¹¹. Legend says that the first version of Git was written in a couple of weeks¹².

CVS

Never heard of CVS? It was a source control system that Joel Spolsky described in September 2000 as *fine* (emphasis his) in the first item of his eponymous Joel Test¹³ for better software:

I’ve used commercial source control packages, and I’ve used CVS, which is free, and let me tell you, CVS is *fine*.

Yes, the first step for better software was (shocker!) using source control software. (As a side note, I started my professional career as a software developer in 1997, and nope, we did not use source control, not even CVS. Yes, you guessed right: we just saved VBScript files locally and uploaded them via FTP. Too bad if we overwrote changes with one another: you only live once. I had to wait until 2002 to use a source control system for the first time, and for the curious among you, it was Rational ClearCase¹⁴.)

Never heard of Joel Spolsky? Well, he was the co-creator of Stack Overflow, which I guess you *have* used at some point in your career. 24 years ago, Joel was one of the first influencers of the burgeoning field of software engineering. Think Kelsey Hightower, but with more controversial views. Or Steve Yegge, but with less controversial views.

Speaking about Stack Overflow, here is an example of the state-of-the-art of source control in 2008. One of the first-ever questions asked on the site, dated September 8th, 2008, asks precisely what version control system to use¹⁵ for a single-developer workflow. (Interestingly, that question was asked precisely at the same time while in the real world, the 2008 financial crisis¹⁶ was breaking havoc. Our industry lives in a bubble, no doubt about that. But I am

⁶https://en.wikipedia.org/wiki/Apache_Subversion

⁷<https://www.cvs.com/>

⁸https://en.wikipedia.org/wiki/Concurrent_Versions_System

⁹<https://deprogrammaticaipsum.com/issue-61-databases/>

¹⁰<https://en.wikipedia.org/wiki/BitKeeper>

¹¹<https://www.wordnik.com/words/git>

¹²<https://www.linuxjournal.com/content/git-origin-story>

¹³<https://www.joelonsoftware.com/2000/08/09/the-joel-test-12-steps-to-better-code/>

¹⁴https://en.wikipedia.org/wiki/Rational_ClearCase

¹⁵<https://stackoverflow.com/questions/49601/is-there-a-barebones-windows-version-control-system-thats-suitable-for-only-one>

¹⁶[https://en.wikipedia.org/wiki/2007%E2%80%932008_financial_crisis#2008_\(September\)](https://en.wikipedia.org/wiki/2007%E2%80%932008_financial_crisis#2008_(September))

digressing, once again.)

I'm trying to find a source control for my own personal use that's as simple as possible. The main feature I need is being able to read/pull a past version of my code. I am the only developer.

The replies to the question consist of a long catalog of pretty much every version control system known to mankind at that point in time.

Windows' Version Control Odyssey

But let us return to the year 2000: a few days before Joel Spolsky published his Joel Test, Mark Lucovsky¹⁷ gave a talk titled "Windows: A Software-Engineering Odyssey"¹⁸ at the 4th USENIX Windows System Symposium¹⁹ in Seattle, Washington. Mr. Lucovsky was a member of the original Windows NT team from 1988 to the mid-2000s. The PowerPoint slides of the talk are still available online²⁰ at the time of this writing, and I seriously recommend you take a look at them.

Because part of the "odyssey" was, you guessed it, source control. On slide 14 you can learn that Windows NT 3.51 used an "internally developed" system... which was "on life support" by the time of Windows 2000:

To keep a machine in synch was a huge chore (1 week to setup, 2 hours per-day to synchronize)

Oops. Not a great way to onboard new team members. Now you know why the Agile Manifesto²¹, published in 2001, was so revolutionary. Thanks to Raymond Chen²², arguably the most important lecturer of Windows history, we learn the name²³ of said internally developed system:

In the early days, Microsoft used a homemade source control system formally called Source Library Manager, but which was generally abbreviated SLM, and pronounced *slime*. It was a simple system that did not support branching.

In slide 24 of Mark Lucovsky's PowerPoint slides, we learn that Microsoft took the decision to migrate the source code of Windows 2000 to something called "Source Depot". Raymond Chen agrees²⁴:

Shortly after Windows 2000 shipped, the Windows source code transitioned to a source control system known as Source Depot, which was an authorized fork of Perforce.

Why Perforce? The choice had to do²⁵ with the gigantic size of Microsoft Windows' source code base:

The justification is perhaps less relevant than it once was, but Perforce tends to perform better on large repositories than Subversion. This is one of the rea-

¹⁷<https://www.linkedin.com/in/mark-lucovsky-5280034/>

¹⁸<https://www.usenix.org/legacy/events/usenix-win2000/invitedtalks.html>

¹⁹<https://www.usenix.org/legacy/events/usenix-win2000/>

²⁰https://www.usenix.org/legacy/events/usenix-win2000/invitedtalks/lucovsky_html/Lucovsky.ppt

²¹<https://agilemanifesto.org/>

²²<https://devblogs.microsoft.com/oldnewthing/>

²³<https://devblogs.microsoft.com/oldnewthing/20180122-00/?p=97855>

²⁴<https://devblogs.microsoft.com/oldnewthing/20180122-00/?p=97855>

²⁵<https://softwareengineering.stackexchange.com/a/85891>

sons Microsoft acquired a source license to Perforce to build Source Depot; NT's repository is a monster, and not many products, commercial or otherwise, could handle it.

Mark Lucovsky's summarized the benefits of Source Depot in two bullet points on slide 24 of his presentation:

- New machine setup 3 hours vs. 1 week
- Normal sync 5 minutes vs. 2 hours

Is the Microsoft Windows team still using Source Depot today? Apparently not. In 2017, we learned that Microsoft migrated all 300 GB of Windows source code to Git in an article on Ars Technica²⁶, which contains another gem describing the "Microsoft odyssey" in source control systems:

Long ago, the company had a thing called SourceSafe, which was reputationally the moral equivalent to tossing all your precious source code in a trash can and then setting it on fire thanks to the system's propensity to corrupt its database.

(I can confirm. Sadly, I should say.)

Microsoft's adoption of Git, however, was not without hurdles²⁷, and led to the creation of the Git Virtual File System (GVFS) project:

But Git isn't designed to handle 300GB repositories made up of 3.5 million files. Microsoft had to embark on a project to customize Git to enable it to handle the company's scale.

The Age Of Git

The infatuation of Microsoft with Git reached its peak in 2018, when it swallowed GitHub²⁸, the platform that arguably made Git mainstream. Three years prior, sensing *l'air du temps*, they had released Visual Studio Code²⁹, with integrated Git support.

GitHub introduced the concept of Pull Requests to the world as early as February 2008³⁰, a feature later adopted and adapted by GitLab³¹, Gitea³² (and its recent fork Forgejo³³), and BitBucket³⁴, and which became the bread-and-butter for code reviews during the past 15 years. But the matter of fact is that GitHub also created a paradox in the world of Git: suddenly, a distributed source control system... became centralized. Some are understandably aghast³⁵ by this state of things.

We are in 2024, and Git is everywhere. The long evolution that led to the Git supremacy in the 2010s and, apparently, also the 2020s, can be summarized as a sequence of open-

²⁶<https://arstechnica.com/information-technology/2017/02/microsoft-hosts-the-windows-source-in-a-monstrous-300gb-git-repository/>

²⁷<https://arstechnica.com/information-technology/2017/05/90-of-windows-devs-now-using-git-creating-1760-windows-builds-per-day/>

²⁸<https://news.microsoft.com/2018/06/04/microsoft-to-acquire-github-for-7-5-billion/>

²⁹https://en.wikipedia.org/wiki/Visual_Studio_Code

³⁰<https://github.blog/2008-02-23-oh yeah-there-s-pull-requests-now/>

³¹https://docs.gitlab.com/ee/user/project/merge_requests/

³²<https://docs.gitea.com/next/usage/pull-request>

³³<https://forgejo.org/docs/latest/user/pull-requests-and-git-flow/>

³⁴<https://support.atlassian.com/bitbucket-cloud/docs/tutorial-learn-about-bitbucket-pull-requests/>

³⁵<https://blog.edwardloveall.com/lets-make-sure-github-doesnt-become-the-only-option>

source programs, one replacing the other: SCCS³⁶ in the 1970s, RCS³⁷ in the 1980s, CVS in the 1990s, and Subversion in the 2000s. To ensure smooth migration paths, Subversion could import CVS repositories, and Git can import Subversion repositories. But most importantly, version control systems migrated from local-only systems like SCCS and RCS, to client-server architectures (CVS and Subversion), to distributed systems, like Git and others, most notably Mercurial³⁸.

(Speaking about Mercurial, did you know that the Firefox developers recently decided to drop it³⁹ and use Git instead?)

These days, we are used to *cloning* an entire project on our computer, after which we can safely plug it off the network and continue writing software in a completely disconnected way. This simple paradigm was utterly and completely unthinkable 20 years ago. And guess what: your local repository also contains the full history of every single change ever made to your project. This was a feature that, naturally, client-server systems could never provide (spoiler alert: the server had the full history, while clients had only the HEAD, so to speak).

Git (and its cheap branching facilities) had a lasting impact on developer workflows. Vincent Driessen published in January 2010 a seminal article titled “A successful Git branching model”⁴⁰ introducing the world to the controversial⁴¹ concept of git-flow. Why controversial? Well, because most opinions in the software industry are such. Now there is a GitHub flow⁴² and an Atlassian Gitflow workflow⁴³ and many more branching workflows available.

Git repositories have become eventful, with every push, merge, or tag operation triggering a workflow somewhere. A whole industry has sprung up, including names such as Argo CD⁴⁴, GitHub Actions⁴⁵, GitLab CI/CD pipelines⁴⁶, and Gitea Runner⁴⁷, providing a new level of automation and convenience. The influence of Git is so strong in this space that the term GitOps⁴⁸ now refers to a whole subset of our industry. But you should not be using branches for deployments, you have been warned⁴⁹.

The question is simple now: what comes after Git? At this point, it is probably impossible to challenge the immense popularity of Git. I say “*probably*” because in our industry, it is impossible to predict the future. There are two interesting contenders worthy of mention: Pijul⁵⁰, written in Rust (although the project started with OCaml⁵¹ a decade ago) or Fossil⁵², written by the creators of SQLite. In this last case, the SQLite team provides a list of reasons⁵³ why not to use Git:

Let’s be real. Few people dispute that Git provides a suboptimal user experience.

³⁶https://en.wikipedia.org/wiki/Source_Code_Control_System

³⁷https://en.wikipedia.org/wiki/Revision_Control_System

³⁸<https://www.mercurial-scm.org/>

³⁹<https://glandium.org/blog/?p=4346>

⁴⁰<https://nvie.com/posts/a-successful-git-branching-model/>

⁴¹<https://web.archive.org/web/20111007075151/http://scottchacon.com/2011/08/31/github-flow.html>

⁴²<https://docs.github.com/en/get-started/using-github/github-flow>

⁴³<https://www.atlassian.com/git/tutorials/comparing-workflows/gitflow-workflow>

⁴⁴<https://argoproj.github.io/cd/>

⁴⁵<https://docs.github.com/en/actions>

⁴⁶<https://docs.gitlab.com/ee/ci/pipelines/>

⁴⁷<https://about.gitea.com/products/runner/>

⁴⁸<https://www.gitops.tech/>

⁴⁹<https://codefresh.io/blog/stop-using-branches-deploying-different-gitops-environments/>

⁵⁰<https://pijul.org/>

⁵¹<https://github.com/8l/pijul>

⁵²<https://fossil-scm.org/>

⁵³<https://sqlite.org/whynotgit.html>

A lot of the underlying implementation shows through into the user interface. The interface is so bad that there is even a parody site that generates fake git man pages⁵⁴.

While we wait for a better alternative, let us read the man 7 giteveryday page and call git-extras⁵⁵, SourceTree⁵⁶, TortoiseGit⁵⁷, Fugitive⁵⁸, Codeberg⁵⁹, and Magit⁶⁰ to the rescue. It seems that, whether we like it or not, we will probably be storing our source code in Git repositories for the next twenty years.

Cover photo by Praveen Thirumurugan⁶¹ on Unsplash⁶².

⁵⁴<https://git-man-page-generator.lokalto.net/>

⁵⁵<https://github.com/tj/git-extras>

⁵⁶<https://www.sourcetreeapp.com/>

⁵⁷<https://tortoisegit.org/>

⁵⁸<https://github.com/tpope/vim-fugitive>

⁵⁹<https://codeberg.org/>

⁶⁰<https://magit.vc/>

⁶¹https://unsplash.com/@praveentcom?utm_content=creditCopyText&utm_medium=referral&utm_source=unsplash

⁶²https://unsplash.com/photos/a-book-and-a-small-figurine-on-a-desk-KPAQpJYzH0Y?utm_content=creditCopyText&utm_medium=referral&utm_source=unsplash

Linus Torvalds

Adrian Kosmaczewski

March 4th, 2024



The tech industry is a fertile ground for anecdotes starred by individuals qualifying as brilliant jerks, psychopaths, and other atrocious types of personality. Suffice to say that, from the height of their positions of supposed leadership, many chose the easy path of advancing a twisted agenda that feeds into their hubris, in the detriment of the wider advancement of society.

Let us name a few such individuals: Larry Ellison¹, Bill Gates², Mark Zuckerberg³, Steve Jobs⁴, Richard Stallman⁵, John McAfee⁶, Peter Molyneux⁷, Travis Kalanick⁸, Adam Neumann⁹, Scott McNealy¹⁰, Elon Musk¹¹, Jeff Bezos¹², Peter Thiel¹³, and Evan Spiegel¹⁴. The list could go on and on, but for the sake of brevity, we will stop here. It should be no surprise to anyone that Linus Torvalds (or at least, his pre-2018 persona) belongs to this group, too.

Just like Reed Hastings¹⁵ and Satya Nadella¹⁶ have decided to get rid of brilliant jerks in their

¹https://bobsutton.typepad.com/my_weblog/2007/03/the_ashholiness.html

²<https://www.businessinsider.com/after-talking-to-teens-last-year-about-anger-management-bill-gates-admits-that-some-of-his-famous-microsoft-tantrums-might-not-have-been-wise-2019-1?r=US&IR=T>

³<https://www.chieflearningofficer.com/2019/02/28/brilliant-jerks/>

⁴<https://www.businessinsider.com/steve-jobs-jerk-2011-10?r=US&IR=T>

⁵<https://thenextweb.com/news/free-software-icon-richard-stallman-has-some-moronic-thoughts-about-pedophilia>

⁶<https://www.wired.co.uk/article/dangerous>

⁷<https://www.somethingawful.com/news/five-biggest-jerks/>

⁸<https://www.forbes.com/sites/richkarlgaard/2014/12/10/do-jerks-always-win/?sh=80e106123251>

⁹<https://www.businessinsider.com/is-your-ceo-brilliant-jerk-or-both-how-to-tell-2019-10?r=US&IR=T>

¹⁰<https://www.zdnet.com/article/scott-mcnealy-has-a-big-mouth/>

¹¹https://www.leadershipnow.com/leadingblog/2020/05/no_brilliant_jerks_how_to_deal.html

¹²<https://sea.mashable.com/tech/15683/how-bezos-became-an-asshole-and-7-other-things-we-learned-from-amazon-unbound>

¹³<https://www.electronicweekly.com/blogs/mannerisms/genius/good-asshole-2014-10/>

¹⁴<https://techcrunch.com/2014/05/28/confirmed-snapchats-evan-spiegel-is-kind-of-an-ass/>

¹⁵<https://www.inc.com/jim-schleckser/why-netflix-doesn-t-tolerate-brilliant-jerks.html>

¹⁶<https://www.businessinsider.com/nadella-brilliant-jerk-phenom-in-tech-is-done-2019-5?op=1&r=US&I>

own corporations, our world at large would benefit from discouraging people from treating others like shit. To a large degree, the types of personality that engage in such behavior do not realize the harm caused around them, in particular, in the long term. In the short term, such words can cause laughter, but they also impregnate the minds of those who attentively listen, and create a tacit sense of approval to engage in similar behavior.

This is a terrible, horrendously vicious circle, and it has to stop.

The Vidéothèque movie of this month¹⁷, published in 2007, is a sad example of this situation. It was one of the Google TechTalks we mentioned last year¹⁸ in this magazine. This one features Linus Torvalds explaining the history and architecture of Git, the most widely used source control system at the time of this publication but an entirely new kid on the block when the video was filmed.

The Good Parts

Let us start with the more benign parts. After diving into the history of the Git, the architectural talk is divided into three sections, each highlighting a particular aspect of the design and architecture of Git: its distributed nature, its high performance by design, and its absolute reliability. Taken together, these three aspects provide a simple answer to the enduring success that Git has had in the past 17 years, certainly explaining its becoming the *de facto* standard for source code management in our industry.

Speaking about its distributed aspect, and apart from the obvious (and dismissive) contrasts with CVS (“tarballs and patches are a much superior source control management system¹⁹ than CVS”) and Subversion, Linus mentions the Monotone²⁰ version control system as an interesting inspiration, but lacking in performance²¹. He also mentions Mercurial²², as having mostly the same design as Git²³. He highlights the importance of being offline²⁴ and how it literally means working in your own branch (we take branching for granted today, but 17 years ago, this was quite a revolutionary concept). Finally, he talks about how Git entered the enterprise space, first used as a client to access Subversion repositories²⁵ (something I have personally done in the past when Subversion was still around). Also, related to corporations, he highlights an important, often overlooked aspect: that in a corporate environment, centralized systems work better²⁶.

Regarding the performance part, there are two significant points. First, that the mistake made by the designers of Subversion was to focus on the wrong problem: branching was not the issue; merging was²⁷. Second, that Git is fast because its code and data structures are really, really simple²⁸.

Finally, Linus mentions Git’s reliability as a result of an actual break-in attempt²⁹ in the

R=T

¹⁷<https://www.youtube.com/watch?v=4XpnKHJAok8>

¹⁸<https://deprogrammaticaipsum.com/google-techtalks/>

¹⁹<https://youtu.be/4XpnKHJAok8?t=171>

²⁰[https://en.wikipedia.org/wiki/Monotone_\(software\)](https://en.wikipedia.org/wiki/Monotone_(software))

²¹<https://youtu.be/4XpnKHJAok8?t=721>

²²<https://www.mercurial-scm.org/>

²³<https://youtu.be/4XpnKHJAok8?t=1942>

²⁴<https://youtu.be/4XpnKHJAok8?t=883>

²⁵<https://youtu.be/4XpnKHJAok8?t=1993>

²⁶<https://youtu.be/4XpnKHJAok8?t=2324>

²⁷<https://youtu.be/4XpnKHJAok8?t=3174>

²⁸<https://youtu.be/4XpnKHJAok8?t=3327>

²⁹<https://youtu.be/4XpnKHJAok8?t=3567>

Linux kernel source code. This resulted in a natural outburst of paranoia, which in turn brought the choice of hashing every single change with SHA-1 (it is worth remembering that this algorithm was considered cryptographically secure³⁰ in 2007, but “SHAttered”³¹ in 2017). Git is not only designed to handle as many files as needed, but also so that developers can trust each one³² of them, at all times. This was yet another revolutionary concept back then.

The Unnecessary Parts

Mixed with the juicy technical bits, there is a third aspect of this talk that makes it relevant, but for all the wrong reasons: a seemingly endless stream of criticisms, negativity, and (supposedly) self-deprecating remarks that end up being punches against anyone or anything Linus happens to dislike.

Let us enumerate some gems: “the designers of Subversion were complete morons³³”; “nobody actually creates perfect code the first time around except me³⁴”; “I think most of you are completely incompetent³⁵”; “the KDE people are... well...”³⁶; “when I merge from somebody, I trust them, but on the other hand, they might have stopped using their medication³⁷”; *und so weiter*.

Yes, there is a “content advisory” at the beginning³⁸ of the talk. But none of this is needed, so let us focus on the good parts, of which there are many.

Balancing Acts

Maybe the issue at the core is that I am getting old. Because yes, I vividly remember laughing loudly when I watched this talk for the first time in 2007, and engaging in similar remarks and ways of thinking. Older entries of my own blog³⁹ are filled with such vitriol. I have given talks (some of which are on YouTube⁴⁰) with such remarks. So, yes, I plead guilty to condoning such behavior, your honor.

I do not approve this anymore. With hindsight, I realize that we (as a society) do not need any more of this. And please refrain from arguing on Reddit or elsewhere that this is just me lacking humor. First, yes, I never had any. And second, it is objectively repugnant, no matter how you look at it. So, you have been warned; definitely watch this month’s Vidéothèque entry⁴¹ of Linus Torvalds, but be aware of foul remarks here and there.

But I have to be fair. 17 years have passed since this video, and in 2018 Linus has apologized⁴² for years of ranting, and I think that as a result, the whole IT field (and in particular, the Linux galaxy) is a far better place because of that.

³⁰<https://youtu.be/4XpnKHJAok8?t=3465>

³¹https://en.wikipedia.org/wiki/SHA-1#SHAttered_%E2%80%93_first_public_collision

³²<https://youtu.be/4XpnKHJAok8?t=3437>

³³<https://youtu.be/4XpnKHJAok8?t=3081>

³⁴<https://youtu.be/4XpnKHJAok8?t=1339>

³⁵<https://youtu.be/4XpnKHJAok8?t=1685>

³⁶<https://youtu.be/4XpnKHJAok8?t=2770>

³⁷<https://youtu.be/4XpnKHJAok8?t=3250>

³⁸<https://youtu.be/4XpnKHJAok8?t=480>

³⁹<https://akos.ma/>

⁴⁰<https://www.youtube.com/@akosma>

⁴¹<https://www.youtube.com/watch?v=4XpnKHJAok8>

⁴²<https://arstechnica.com/gadgets/2018/09/linus-torvalds-apologizes-for-years-of-being-a-jerk-takes-time-off-to-learn-empathy/>

Even more: I am positively in awe when faced with the incomparable talent of Linus and the brilliancy of his creations. His work has been a definitive positive source of change and wealth for the whole planet; mankind is *collectively better* because of his work. The Linux kernel *and* the Git version control system are both incommensurately valuable and brilliant, and have had (and will have) a transforming impact on the history of mankind. Each in itself would have been a remarkable achievement; taken together, they confirm the magnificence of the creative mind at their origin.

(There is also a possible argument to be made here, that Linux and Git are so good *precisely* because Linus cares so much, and is so passionate about his creations. I do not know him in person, but I think there might be a certain truth in this possibility, too. I certainly hope so.)


However brilliant, we do not need any more brilliant jerks spewing abusive thoughts towards other creators, particularly when their only sin was that of coming up with ideas that you just happen to dislike. We do not need such statements neither on record nor in other public or private spaces, no matter how good they are, or how big or small their impact on this planet may be. We should collectively let our work speak for ourselves, direct our cynicism and anger toward the resolution of actual world problems (there are a lot of those to choose from), and remember to have some respect for one another.

Cover snapshot chosen by the author from the video.

Travis Swicegood

Adrian Kosmaczewski

March 4th, 2024



A quick review of previous entries in the Library section of this magazine shows that it does not feature any book from The Pragmatic Programmers¹, for no other reason than gross oversight. We have discussed books from MIT Press, Addison-Wesley, O’Reilly, and many other publishing houses, and now it is time to solve this issue. This month we will elaborate on “Pragmatic Version Control Using Git”², a 2008 book by Travis Swicegood³.

In other circles than software, such as physics or literature, the year 2008 would undoubtedly be called an *annus mirabilis*. While the “real world” was grappling with one of the most severe financial crisis⁴ in history, the software industry saw a miraculous conjunction of factors, seemingly happening within months of one another. Let us enumerate a few: in March, the availability of the iPhone SDK. In April, the launch of GitHub. And in September, three major events: the announcement of the Google Chrome web browser on the 2nd, the launch of Stack Overflow on the 15th, and the release of Android 1.0 on the 23rd.

Couple all this with the meteoric rise in popularity of AWS and cloud computing in general, the phenomenon of social media, and the explosion of Web 2.0 and “Ajax”⁵, with the ubiq-

¹<https://pragprog.com/>

²<https://pragprog.com/titles/tsgit/pragmatic-version-control-using-git/>

³<https://github.com/tswicegood>

⁴https://en.wikipedia.org/wiki/2007-2008_financial_crisis

⁵[https://en.wikipedia.org/wiki/Ajax_\(programming\)](https://en.wikipedia.org/wiki/Ajax_(programming))

uitous presence of libraries such as jQuery⁶ and Prototype.js⁷. You could feel that software engineering was never going to be the same from that point onwards. It was undergoing a profound transformation, live, in front of our eyes.

It was both a terrific and terrible time to be a software developer, and Git appeared as one of the hottest new tools to check out. But Git was also substantially different from anything many of us had used up to that point. At least for me, the whole idea of a distributed source code management system was completely alien.

Up to that moment I had been a happy Subversion⁸ user, hosting some of my code on Google Code⁹ and even running Subversion repository servers on my own machine—I admit, a rather clumsy setup for a single developer. Before Subversion, I had used various source code management tools: Rational ClearCase¹⁰ from 2002 to 2003, Visual SourceSafe¹¹ from 2004 to 2005, and CVS¹² from 2006 to 2007. (Yes, the list shows a clear regression; it seemed that the further I advanced in my career, the dumber and more unstable the SCM of my employer would be. *Anyway.*)

All the systems enumerated above are client-server, and thus, centralized; there is a server, somewhere, and a small client on my laptop allows me to check code out, commit changes, and voilà. Need to blame a file? Talk to the server. Want to see the full history log? Talk to the server. Want to create a branch? You get the gist.

I remember attending a talk about Git during a community meetup in March 2008, and rushing back home to order a book about it on Amazon (Kids: e-books were just starting to appear on our radar, as the original Kindle¹³ was launched in November 2007. So, physical book it was.) There were not many books published about Git (yet) and I had always enjoyed reading books from The Pragmatic Programmers, so choosing this month's Library book was a no-brainer for me.

Needless to say, this was the perfect book to open my mind to the possibilities of Git. Just like in many other Pragmatic Bookshelf titles, the introduction to the subjects is gentle, complete, yet entertaining. The first chapter was particularly helpful, with its explanation of the differences between distributed and centralized version control, and the resulting workflows that were possible in each case.

The conceptual part was the most challenging one for me, but the author made it abundantly clear that it was not *that* hard. And the highly practical nature of the book, filled with examples, made understanding Git fast, straightforward, and simple. The detachable “Git Command Quick Reference” sat on my desk for months, providing a much-needed reference until most commands were engraved in my mind.

The rest is history. When I started my own business¹⁴ in 2008, I got a paying GitHub account and used it to host the private Git repositories of all of my customers. But some of those customers were not entirely used to Git or distributed source code management for

⁶<https://en.wikipedia.org/wiki/JQuery>

⁷<http://prototypejs.org/>

⁸<https://subversion.apache.org/>

⁹<https://code.google.com/>

¹⁰https://en.wikipedia.org/wiki/Rational_ClearCase

¹¹https://en.wikipedia.org/wiki/Microsoft_Visual_SourceSafe

¹²https://en.wikipedia.org/wiki/Concurrent_Versions_System

¹³https://en.wikipedia.org/wiki/Amazon_Kindle

¹⁴<https://akos.ma/blog/running-akosma-software/>

that matter, and asked me to push code into their Subversion repositories; no problem at all, `git svn` to the rescue¹⁵.

Git was a godsend during those years of independent consulting. I was very frequently on the go, either on a plane or on a train, without any network connectivity, and Git allowed me to work perfectly well, committing, diff'ing, blaming (mostly me) and branching as much as I needed, at any time, in any place. As soon as a network connection became available, a simple `git push` would share those changes with my customers, and even better, trigger some automated task in a CI/CD system somewhere.

Beyond my own developer experience, one thing that changed in the software industry from 2008 onwards was that Git became the all-encompassing source code management tool, while others slowly disappeared into oblivion. New developers joining the workforce learned Git exclusively. New businesses started adopting Git and collaborating on GitHub or GitLab repositories. Popular software projects started opening up shop on GitHub. Text editors started providing built-in Git support. Everyone started assuming that Git was the big default SCM of the world of software. Well-known SCM packages slowly faded away. Companies providing commercial SCMs went bankrupt or pivoted to other markets.

During the past 16 years, I almost exclusively used Git for my code and my projects, both personal and professional. I did use Mercurial too, for a while at least, but for no other reason than my employer at the time was a heavy Python¹⁶ user. They ended up migrating to Git at some point, the power of networking effects being too difficult to counter. Those forces even bent BitBucket¹⁷, originally a Django web application hosting only Mercurial repositories, to drop Mercurial support in 2020. *Sign o' the times*.

The writing was on the wall. Git is the *de facto* standard, and it has become as ubiquitous as the English language, or water. Travis Swicegood's book still sits on my bookshelf; admittedly, I do not consult it so often as I did in those early days, but it represents a pivotal moment for my career—and truly, for the whole industry.

Cover photo by the author.

¹⁵<https://git-scm.com/docs/git-svn>

¹⁶<https://deprogrammaticaipsum.com/the-state-of-python-in-2021/>

¹⁷<https://en.wikipedia.org/wiki/Bitbucket>