

Professionalism on the LKML: Metadiscourse about Open Source Development

Liz Fong-Jones, MIT course 6-3 student (class of 2014).

Written in November 2013 for 21A.503: Language and Technology.

This paper was originally written in 2013; Sage now uses they/them pronouns, but source material of mailing list posts still shows their previous name/pronouns. In March of 2015, a "Code of Conflict" process document was added to the kernel tree. Linus went on to finally [apologize for his behavior](#) and [accept addition of a Code of Conduct](#) modeled after the [Contributor Covenant](#) in September 2018, 5 years later, in response to pressure from a [New Yorker article](#) with similar analysis as this paper.

Over the past decade, open source development has changed from a small, niche hobbyist community to a sprawling bazaar comprised of tens of thousands of individuals. Participants in open source rely upon electronic communication tools such as mailing lists and Internet Relay Chat (IRC) to mediate interactions between geographically distributed developers and ensure the flow of technical discourse about project development. Although developers occasionally collaborate face to face at physical conferences, the day to day work and the primary locus of the development community remain on mailing lists. One prominent open source project, the Linux kernel, was first established in 1991 as a "just a hobby, [not] big and professional like gnu" by a single author, Linus Torvalds. Today, it involves thousands of contributors per release, including employees from hundreds of for-profit companies who are paid to make their contributions¹.

¹ 1392 developers from 243 distinct companies made contributions to the most recent kernel release; 9784 developers from 1064 companies have contributed to the Linux kernel since 2005. Linux Foundation. "Linux Kernel Development: How Fast It is Going, Who is Doing It, What They Are Doing and Who is Sponsoring It" (2013) pp. 7.

With such a large and globally distributed community that cannot convene in the same physical space at the same time, Linux kernel developers primarily interact on the Linux Kernel Mailing List (LKML) email list. Discussions about the kernel development process and proposed patches are pushed in realtime to all approximately 7000 subscribers (as of October 2013), and are also publicly archived to provide a historical record for accountability and transparency. Thus, researchers can easily study the community through its archive to understand the discourses about software development and metadiscourses about the list itself that take place there. One such recent trend of metadiscourse involves the increasing involvement of companies and their employees in Linux development, introducing a tension into the development process between more tenured contributors from the hobbyist era who seek to continue their previous brusque styles of leadership and employees who view the community as an extension of their workplace in which civility should be a norm. This tension is resolved through metadiscourse about permissible discourse on the lists in order to reconcile the media ideologies of the community's members, and ultimately through decisions to media switch to defuse disagreements.

The tensions between this pair of conflicting media ideologies came to a head on the LKML on July 15, 2013. Sage Sharp, an employee of Intel and one of the sub-maintainers of the Universal Serial Bus subsystem, interjected their concerns about civility into a thread in which their upstream subsystem maintainers were participating, and in which Linus openly encouraged his lieutenants to enforce standards by making verbal threats. Linus stated his belief that such techniques were effective for reducing

errors introduced into the development process. Sage argued that Linus's techniques were counterproductive, and would lead to individuals feeling intimidated and leaving the community rather than be subject to harassment and threats.

Subject: Re: [00/19] 3.10.1-stable review

Seriously, guys? Is this what we need in order to get improve -stable? Linus Torvalds is advocating for physical intimidation and violence. Ingo Molnar and Linus are advocating for verbal abuse.

*Not *fucking* cool. Violence, whether it be physical intimidation, verbal threats or verbal abuse is not acceptable. Keep it professional on the mailing lists.*

Let's discuss this at Kernel Summit where we can at least yell at each other in person.

Yeah, just try yelling at me about this. I'll roar right back, louder, for all the people who lose their voice when they get yelled at by top maintainers. I won't be the nice girl anymore.²

Subject: Re: [00/19] 3.10.1-stable review

"No, I can't opt-out. I'm paid to work on the Linux kernel. 80% of the Linux kernel developers are paid to work on Linux. This is not about a volunteer project anymore. This is about a person paid to manage the kernel (Linus) verbally abusing other kernel developers who are also paid to work on Linux. This is about making a workplace (the Linux kernel community) a safe workplace, even during heated technical discussions."³

In this discussion, they raised the question of whether Linux constitutes merely a hobby community in which anything goes, or whether it is a workplace with all the associated norms associated with a conventional workplace. Their initial set of posts spawned hundreds of comments on their blog, approximately one thousand replies on the LKML, and hundreds of comments on technology news sites as they covered the flame war. The sheer length of this metadiscussion in-band on the LKML was noteworthy because it broke from the norms of the LKML to focus primarily on technical issues rather than upon metadiscourse.

² Sage Sharp. "Re: [00/19] 3.10.1-stable review".

<http://marc.info/?l=linux-kernel&m=137390362508794&w=2>

³ Sage Sharp. Comment on "No More Verbal Abuse".

<http://sage.thesharps.us/2013/07/15/no-more-verbal-abuse/#comment-1749>

In order to examine this tension surrounding language and behavior on the LKML, it is necessary to first understand the positions of both sides of the debate. The incumbent group consists of the senior kernel developers responsible for initially growing the community during its hobbyist phase. Linus Torvalds has led the development of the official Linux kernel since its inception in 1991, and set the initial tone for discourse about kernel development. In his role as "benevolent dictator for life"⁴ due to having founded the project and being entrusted by the community with its stewardship, he holds veto power over what patches authored are pulled into the canonical version of the kernel. His communication style emphasizes uncompromising direct feedback and strict enforcement of his "key rules" for the kernel that must never be broken, with the goal of encouraging high standards from his immediate subordinates and thus, indirectly from the entire community. His stated norms are to "be polite when it's called for... [but] when people who know better send me crap, I'll curse at them... I expect more from [top developers], and conversely I'll be a lot more upset when they do something that I really think was not great."⁵ Rather than directly micromanage merging each commit, Linus delegates responsibility for each major subsystem to subsystem maintainers, nicknamed 'lieutenants', and those subsystem maintainers accept relevant merge requests from other developers, or from subsubsystem maintainers. When his trusted lieutenants make egregious mistakes, he chastises them in public on the LKML. Thus, Linus's media ideology for mailing lists that it is acceptable to swear at contributors in an attempt to

⁴ Steve Hamm. "Linus Torvalds' Benevolent Dictatorship". Businessweek (2004).
<http://www.businessweek.com/stories/2004-08-17/linus-torvalds-benevolent-dictatorship>

⁵ Linus Torvalds. "Re: [00/19] 3.10.1-stable review".
<http://marc.info/?l=linux-kernel&m=137391528412980&w=2>

discourage subpar work propagates to the community hierarchically through his lieutenants to contributors to each component.

In contrast to the relatively freewheeling volunteerism of the open source community, software developers working for companies have a different set of expectations about their working environment and the media ideology of communications appropriate for public mailing lists. They expect to treat others with professionalism, and in turn be treated civilly by their peers and managers. Personal attacks are forbidden and human resources departments can intervene if the community's expectations of behavior are violated. Criticism is typically delivered in less strong terms and in a private one on one setting rather than in public on mailing lists. Therefore, corporate developers encounter frustration if they must subject themselves to what they perceive as 'harassment' or 'bullying' in the course of their jobs to interact with the open source community. However, among kernel developers there is considerable dispute as to whether behaving in a more 'professional' manner merely means "no bullying" and mutual respect, or instead creates tolerance for mediocrity and a facade of doubletalk and office politics. One developer, Neil Brown, positively characterized employing civility even when discussing poorly-received code by stating: "Being 'polite' without being 'nice' is quite possible. It even has a name: Diplomacy."; however, Joe Perches opposed diplomatic conversations as "circular/indirect/implied/useless," exemplifying Linus's attitude⁶.

⁶ Joe Perches and Neil Brown. "Re: [00/19] 3.10.1-stable review".
<http://marc.info/?l=linux-kernel&m=137393226318188&w=2>

The heated discussion on the LKML eventually wound down due to the efforts of multiple individuals to pull the discussion off-list, including both Sage Sharp and Linus Torvalds, the original flash points for the debate, and other respected third parties such as Theodore Ts'o, a prominent filesystems subsystem maintainer. Parallels can be drawn between the media ideology of best practices for conflict resolution in open source culture by media-switching and the media ideology of individuals changing media while negotiating breakups studied by Gershon in 2010⁷. "Take it offline" is a common expression used whenever discussions become too heated, under the supposition that disputes are better managed face to face, or at minimum over voice rather than over a mailing list or text-only chat room. Sage Sharp's mailing list messages explicitly hearken to this, with their stated desire to resolve the issues at the physical Kernel Summit conference in order to avoid their perceived challenges with using mailing lists to make their point. Rather than "moving between two media chang[ing] the formal nature of the exchanges one is having" (Gershon 2010: 394), communications on-list and off-list are equally informal, and the difference is instead the amount of non-verbal cues available to make intent clear and defuse arguments. Theodore Ts'o implored the LKML community to allow further discussions to take place at an already planned conference rather than continue to hash out the issue in-band on the mailing list; his message also changed the subject of the thread, signaling a clear shift in topic and desire for his message to stand alone as a top-level response rather than as simply one of hundreds of replies to Sage's original post.

⁷ Gershon, Ilana. "Breaking Up Is Hard To Do: Media Switching and Media Ideologies". *Journal of Linguistic Anthropology*, Vol. 20, Issue 2, pp. 389–405 (2010).

Subject: Maybe it's time to shut this thread down (Was: Re: [00/19] 3.10.1-stable review)

The reason why I started the kernel summit over ten years ago was because there were certain topics that are much better discussed in person, and that over time, if we don't have sufficient face to face interactions, the quality of e-mail discussions can start to become frayed.

*One of the reasons is that e-mail is just not as expressive a medium as face-to-face conversations. As a result, when people feel that they aren't being heard, because they aren't getting those critical non-verbal cues, they start escalating. They start using stronger words, such as F*CK. They start doing exactly what they claim to abhor to their verbal opponents in the debate, which is describing their fellow kernel developers using demeaning terms. They start using loaded, and over-reaching words, like "abuse", which ultimately ends up hurting their own case.*

I suspect this is happening because it's easy when a body feels that their message of say, "could we please treat each other with more respect", isn't getting heard, it's very easy and very tempting to resort to "Linus is an AB-UUUUUUUU-SER!".

May I make the polite suggestion (and we'll see how well polite requests get honored via e-mail), that we take this discussion off-line, and wait to try to discuss this in person at the Kernel Summit?⁸

In closing, Sage stated, "I'm not trying to shut down this discussion. But please, let's continue this discussion at KS, away from the court of public opinion. I would love for this email to serve as a final summary of my opinion. We can use this email to start a conversation at KS, and we can argue our hearts out there about the various points."⁹

Linus echoed, "Ok, everybody back to work... Everybody has been gossiping around the water cooler instead, and I've had one or two pull requests per day from the people who are happily oblivious (or too smarter to get involved) about the whole flame fest...

clearly all the energy has been going into arguing, not working. Shoo, shoo. Back to your

⁸ Theodore Ts'o. "Maybe it's time to shut this thread down (Was: Re: [00/19] 3.10.1-stable review)". <http://marc.info/?l=linux-kernel&m=137415423423858&w=2>

⁹ Sage Sharp. "Re: mistakes in code vs. maintainer flow mistakes (was: [00/19] 3.10.1-stable review)". <https://lkml.org/lkml/2013/7/19/634>

desks. Don't all congregate in the break room."¹⁰; in this comment, he explicitly hearkens to the analogy with a traditional workplace, backing down from his previous position that Linux is a hobby community distinct from a workplace.

By the end of the discussion, the parties involved had come to a somewhat better understanding and acceptance of each other's' positions. Sage acknowledged they could not change the behavior of other developers if they accepted attacks from Linus directed at his lieutenants as a norm. On the opposite side, the discussion had generated agreement that such verbal attacks were an exception to which Linus's immediate lieutenants had consented when accepting their roles as leaders, but that such attacks should only be permitted in the context of established relationships. This consensus enabled Sage to focus on communicating to new developers that Linus's behavior towards his trusted, experienced lieutenants would not trickle down to newbies receiving personal attacks for mistakes committed while learning the technical standards. Through the process of metadiscourse, the Linux Kernel Mailing List community brought into the open the elephant in the room about the tension between two different media ideologies, and attempted to bridge the gap in expectations of behavior on the list.

¹⁰ Linus Torvalds. "Ok, everybody back to work..". <https://lkml.org/lkml/2013/7/18/426>

<http://marc.info/?l=linux-kernel&m=137398505801644&w=2>
<http://marc.info/?l=linux-kernel&m=137398593001977&w=2>
<http://marc.info/?l=linux-kernel&m=137398683002304&w=2>
<http://marc.info/?l=linux-kernel&m=137398741802530&w=2>
<http://marc.info/?l=linux-kernel&m=137399874006836&w=2>
<http://marc.info/?l=linux-kernel&m=137400895110430&w=2>
<http://marc.info/?l=linux-kernel&m=137400984610683&w=2>
<http://marc.info/?l=linux-kernel&m=137401007410744&w=2>
<http://marc.info/?l=linux-kernel&m=137401130111098&w=2>
<http://marc.info/?l=linux-kernel&m=137401277611625&w=2>
<http://marc.info/?l=linux-kernel&m=137402304714211&w=2>
<http://marc.info/?l=linux-kernel&m=137402507714691&w=2>
<http://marc.info/?l=linux-kernel&m=137403170916180&w=2>
<http://marc.info/?l=linux-kernel&m=137410815810454&w=2>
<http://marc.info/?l=linux-kernel&m=137411926713140&w=2>
<http://marc.info/?l=linux-kernel&m=137414397820410&w=2>
<http://marc.info/?l=linux-kernel&m=137416255827078&w=2>
<http://marc.info/?l=linux-kernel&m=137416368527483&w=2>
<http://marc.info/?l=linux-kernel&m=137416916529486&w=2>
<http://marc.info/?l=linux-kernel&m=137423224615759&w=2>
<http://marc.info/?l=linux-kernel&m=137441920627214&w=2>
<http://marc.info/?l=linux-kernel&m=137391401112514&w=2>
<http://marc.info/?l=linux-kernel&m=137391528412980&w=2>
<http://marc.info/?l=linux-kernel&m=137391587513153&w=2>
<http://marc.info/?l=linux-kernel&m=137391619513236&w=2>
<http://marc.info/?l=linux-kernel&m=137391800413818&w=2>
<http://marc.info/?l=linux-kernel&m=137392092114772&w=2>
<http://marc.info/?l=linux-kernel&m=137392217115155&w=2>
<http://marc.info/?l=linux-kernel&m=137392506516022&w=2>
<http://marc.info/?l=linux-kernel&m=137392806116971&w=2>
<http://marc.info/?l=linux-kernel&m=137390362508794&w=2>
<http://marc.info/?l=linux-kernel&m=137390810310498&w=2>
<http://marc.info/?l=linux-kernel&m=137390965511100&w=2>
<http://marc.info/?l=linux-kernel&m=137391145411685&w=2>
<http://marc.info/?l=linux-kernel&m=137391223711946&w=2>
<http://sage.thesharps.us/2013/07/15/no-more-verbal-abuse/#comments>
<http://marc.info/?l=linux-kernel&m=137422823914376&w=2>
<http://marc.info/?l=linux-kernel&m=137426050324915&w=2>
<http://marc.info/?l=linux-kernel&m=137449108112389&w=2>
<http://marc.info/?l=linux-kernel&m=137400952810588&w=2>
<http://marc.info/?l=linux-kernel&m=137401282111639&w=2>

<http://marc.info/?l=linux-kernel&m=137401467212190&w=2>
<http://marc.info/?l=linux-kernel&m=137401634712690&w=2>
<http://marc.info/?l=linux-kernel&m=137401863913257&w=2>
<http://marc.info/?l=linux-kernel&m=137402412614476&w=2>
<http://marc.info/?l=linux-kernel&m=137402492014665&w=2>
<http://marc.info/?l=linux-kernel&m=137406622827453&w=2>
<http://marc.info/?l=linux-kernel&m=137407255630179&w=2>
<http://marc.info/?l=linux-kernel&m=137407378430732&w=2>
<http://marc.info/?l=linux-kernel&m=137408047001271&w=2>
<http://marc.info/?l=linux-kernel&m=137408217201857&w=2>
<http://marc.info/?l=linux-kernel&m=137408710803514&w=2>
<http://marc.info/?l=linux-kernel&m=137432377305970&w=2>
<http://marc.info/?l=linux-kernel&m=137407205629947&w=2>
<http://marc.info/?l=linux-kernel&m=137406967728875&w=2>
<http://marc.info/?l=linux-kernel&m=137405586823617&w=2>
<http://marc.info/?l=linux-kernel&m=137401538112420&w=2>
<http://marc.info/?l=linux-kernel&m=137399141104272&w=2>
<http://marc.info/?l=linux-kernel&m=137398978703574&w=2>
<http://marc.info/?l=linux-kernel&m=137394582521250&w=2>
<http://marc.info/?l=linux-kernel&m=137401440812107&w=2>
<http://marc.info/?l=linux-kernel&m=137401314411729&w=2>
<http://marc.info/?l=linux-kernel&m=137400916710490&w=2>
<http://marc.info/?l=linux-kernel&m=137400505408992&w=2>
<http://marc.info/?l=linux-kernel&m=137396693026996&w=2>
<http://marc.info/?l=linux-kernel&m=137396002224746&w=2>
<http://marc.info/?l=linux-kernel&m=137394794821667&w=2>
<http://marc.info/?l=linux-kernel&m=137394454920951&w=2>
<http://marc.info/?l=linux-kernel&m=137393655919224&w=2>
<http://marc.info/?l=linux-kernel&m=137393218218155&w=2>
<http://marc.info/?l=linux-kernel&m=137393016517578&w=2>
<http://marc.info/?l=linux-kernel&m=137392793716941&w=2>
<http://marc.info/?l=linux-kernel&m=137392778416890&w=2>