Three years ago, SIGCOMM began a process of change, including a call for community proposals that received 26 submissions. We adopted several proposals that made changes "at the conference" and some "after the conference," but substantially changing aspects of SIGCOMM "before the conference" (i.e. the paper review process) has taken longer because of the importance and scale of this process. "SIGCOMM beta" is an effort to experiment more practically, by peeling off a fraction of the paper submission traffic and program into an experimental pathway that runs adjacent to the mainline SIGCOMM process. Below is an example proposal for a "SIGCOMM beta" experiment.

Example proposal for a "SIGCOMM beta" experiment: the alternative review pathway (ARP)

Keith Winstein (keithw@cs.stanford.edu) May 1, 2026

Overview: This "SIGCOMM beta" proposal would create an experimental pathway to publishing a SIGCOMM-sponsored conference paper, starting at the SIGCOMM 2028 conference. Authors could submit any time, would receive guaranteed-latency responses and high-quality intellectual engagement with their work, and could choose to iterate with a reviewer until the reviewer "gets it." Reviewers would be able to engage deeply with one paper at a time, ask clarifying questions and focus on papers they feel qualified to review, and include their research groups in the deliberative process for the educational and community-building benefits. At the conference, attendees would benefit from the presence of senior members of the community giving interesting talks that put the papers they champion in context. ARP papers would have talks in the normal SIGCOMM program and would be published in the same proceedings, branded as coming via ARP (they'd be "real" SIGCOMM papers, in order to attract a similar catchment of submissions). *This is a radical proposal*, and I'll be relieved if it isn't selected as the "SIGCOMM beta" experiment because it would be a lot of work to put together. But if I think about what I would do if I were serious about trying to improve the review process and the conference experience in the face of current headwinds, it's something like this.

Goals: For authors, ARP would provide rapid, high-quality intellectual engagement with their work and feedback that authors value, in the sense of improving the work and increasing its impact. For readers, ARP will publish papers where somebody has engaged deeply with it and believes and can explain why reading it will teach the reader something of value. For attendees, we want to increase attendance from senior members of the SIGCOMM community and spur more interesting sessions at SIGCOMM-sponsored conferences.

Submission mechanism: The submission process runs year-round and continuously and is aimed at providing high-quality feedback on a predictable timescale. We will invite past members of the SIGCOMM program committees to serve as "ARP editors." Every editor will have an open, public queue of paper review requests. Review is single-blind in the unconventional direction: the authors are anonymous but the reviewer is not. Everybody on the Internet can see the queue for each editor, including the PDFs of the anonymized paper submissions. Authors will be able to monitor the progress of their paper in the queue and may continue to revise it until it's at the head of the queue. Any author can submit a paper to the back of any queue, with the restrictions that:

- authors and prospective editors will be separated by the standard SIG conflict policy, and
- every author can have only one outstanding submission at any given time, across the whole system.

Admission criteria for each queue (including role of LLMs): Editors may post "admission criteria" for their queue, which can include things like "I'm interested and qualified to review papers on *xyz* topics" or "if it's a system, I like to see a candid limitations section, both about limitations of the proposed system and of the experiments characterizing it" or "before submitting a paper to my queue, please prepend the following prompt and submit as a query to [an LLM chat service available on reasonably priced and nondiscriminatory terms that I don't have a business relationship with], then include a public link to the chat exchange."

Review mechanism: By default, each editor commits that if their queue is nonempty as of Monday morning at 9 a.m. local time, they'll provide a "reply" (decrementing their queue length by 1) before the Monday two weeks later. Editors can select a different reply-rate commitment, e.g. one reply per week, or one reply every three weeks, and it will apply to any subsequently enqueued papers. So based on the number of entries in each queue, the system can promise a submitting author exactly by when they'll get a reply (of course editors are welcome to go quicker through their queues). The reply represents the editor's own independent thought and judgment, as a fiduciary for the eventual reader. In forming their views, it's acceptable for the editor to include unconflicted colleagues, students, or an LLM, but all of these should be credited in the reply and can't substitute for the editor's own independent thought and judgment.

The reply to authors: The "reply" is also public and can be one of four kinds:

1. "this isn't for me—please try somebody else [possibly with suggested names or publication venues]"
2. a standard review, including "for me to champion this, I'd need to see [could include additional experiments, different presentation, exposition in some area, a concrete benefit not yet shown, . . .]"
3. "I'm ready to champion this as long as you make the following shepherding-level changes" (at this point, no more experiments can be requested)
4. "I am going to champion this for publication at SIGCOMM" [in subsequent years, we'd aim to add other target venues, e.g. IMC and CoNEXT]

Editors must not take steps to learn who the authors are—we want papers from anywhere to be on equal footing. In a reply, editors will describe how uncertain they are about the authors' identities. If they think they're pretty sure who the authors are, that's okay but they have to disclose this to keep everybody honest and let us understand how well the system is working.

Championing: If an editor commits to championing the paper, that means:

- they write a 1–3 page commentary endorsing/championing the paper and describing the value it brings readers, with any caveats or context they feel beneficial to a reader's understanding (this is a public, authored document that will be published in the proceedings),
- the commentary can be a co-authored product (e.g. including the champion's research group or colleagues as named co-authors)—it's meant to be an interesting document and should credit whoever contributed,
- the champion presents the paper at an upcoming every-two-months Zoom call of the "normal" SIGCOMM PC (or whatever members of the "normal" PC choose to join the call). The champion advocates for the paper's inclusion while candidly discussing its weaknesses. The PC's involvement is meant to be a lightweight process but also a safety value against giving any individual editor too much power,
- if the "normal" PC accepts the endorsement, then the PC chooses a shepherd (or the champion is the shepherd) who helps the authors prepare it for final publication; the paper will be published at the next SIGCOMM (flagged as an ARP paper, but still in the main SIGCOMM proceedings),
- the authors present the paper at the SIGCOMM conference, in the topic-appropriate session, but also
- the champion commits to attend SIGCOMM in person and deliver a high-quality 5-minute talk about the paper, or at least inspired by the paper, to be given in a dedicated session (the talk can also include contributions from a colleague or advisee who contributed to the endorsement, but the champion is ultimately responsible for the quality of the talk). This can be similar to a "papers we love"-type talk or a broader discussion of the paper's role in context (similar to a "topic preview session") or whatever is most edifying to the audience. Part of the goal is to make champions have "skin in the game," but also, I'm optimistic these talks will be interesting to watch and learn from.

What success would look like: Authors start to prefer submitting to ARP over the normal SIGCOMM pathway, reviewers prefer reviewing for ARP, the ARP-track SIGCOMM publications are of similar quality to the normal track, and the ARP speakers enliven the conference.

Risks (or how this might fail): We're giving the editors a lot of power (not quite the ability to single-handledly "whiteball" a paper, but close); there's a risk of them applying this power unfairly or sloppily, and of more egregious forms of misconduct, e.g. review rings / intentional unblinding. Letting authors choose their editor might be a terrible idea, e.g. maybe some editors will develop a reputation as "easy acceptors" and attract a year-long queue backlog. We might find ourselves needing to limit the number of papers an editor can champion, although the intrinsic costs of championing a paper (writing up an endorsement, presenting to the PC and winning their approval, and preparing a high-quality talk), as well as the queue sojourn delay itself, might provide enough natural backpressure. There's also a big risk that in making the submissions, the reviews, and the reviewer identities public, that both authors and editors will pull their punches. It's not clear that people will want to be editors in this system (maybe they'd be more comfortable if the "review" is a product of a whole research group or of multiple collaborating editors?). We'll need to judge how well this process is going and adjust over time. We're also creating extra work for the "normal" SIGCOMM PC because they'd have to commit to staying on duty to review endorsements until the next PC is appointed.

Schedule: Probably this would be planned nominally as a three-year experiment, run by me and supervised by the SIG Board or TSC. We're not going to get it perfect in the first year so it probably needs to take some time to evolve, similar to how NSDI did "multiple deadlines" as a three-year experiment before making it public. If this is the chosen proposal, we'd probably start recruiting editors and developing the review website/system in October 2026, and we'd want the review system to be online by February or March 2027 (after the SIGCOMM 2027 deadline). Some portion of the SIGCOMM 2027 PC would have to stay on until the next PC is appointed. The first ARP papers would be presented at SIGCOMM 2028. Probably any ARP paper accepted after the SIGCOMM 2028 acceptances (April 2028) will have to wait until SIGCOMM 2029 to be presented because we'd need to hit the same camera-ready deadline as the normal SIGCOMM conference.

Budget: This would be a lot of effort, creating a new review system, recruiting editors, urging authors to try submitting to ARP, and adjusting over time based on what we learn. I'm quite nervous I'd do all this effort and then have to deal with a bunch of LLM-fueled misbehavior (both from authors and editors), or the 100 other ways this could fail (e.g. nobody wants to submit, nobody wants to champion, some misbehavior we haven't accounted for, etc.). I will be relieved when this proposal is not selected! But if I'm serious about what I think is necessary to preserve the joy and intellectual benefit of the conference system, it's something like this. If I end up being chosen do this, I would probably need the SIG to fund the creation of the new reviewing platform and would be asking the SIG to fund part of my summer salary for this.

Why I think this proposal would be good: I just spent 2.5 years as the SIG's "change czar," a position created after the 2023 business meeting and upswell of community proposals aimed at improving various parts of the conference system. My job was basically to nudge and assist the various organs of SIGCOMM towards change; we divided the community-submitted ideas into three categories: "after the conference" (e.g. things like published reviews and ensuring open access), "during the conference" (e.g. the open call for non-paper sessions, and going multi-track to accommodate more papers), and "before the conference," e.g. changes in how reviewers are selected or how reviews are done. This last category is the stickiest because the stakes are arguably the highest. I worked unsuccessfully to broker a deal where NINeS would have been one day of SIGCOMM 2026 (this was the origin of the "SIGCOMM beta" idea), and ultimately everybody got aligned with a plan where a NINeS-adjacent team will chair all of SIGCOMM 2027. This will probably be the biggest "before the conference" change that SIGCOMM has seen in a while, and I hope some of it sticks.

For the longer-term future, my time on the SIG board made me pretty concerned. We face a lot of headwinds, from the LLM-aided misbehavior of authors and reviewers (which consumes a huge amount of time for those unfortunate to have to deal with it), from time-to-publish cycles getting longer through multiple submissions, from academic precarity and institutional focuses on bibliometrics and what that does to the papers, from the world getting bigger without our processes scaling to match, and from AI and other threats to scientific publishing as a means for advancing human intelligence, creativity, and the human condition. I think the ship of SIGCOMM is too big to steer all at once—the annual PC chairs are managing a >100-person committee and hundreds of submissions, and it's hard to give adequate effort to innovations that aren't on the critical path. I'm optimistic that peeling off some of the submissions into a "SIGCOMM beta" pathway will let us experiment more nimbly with the "before the conference" part of things.

I expect ARP will prove too radical for our community to be that experimental pathway, but I think it would improve many of the things that frustrate and concern me about our research conferences today:

- As an author and adviser, I don't enjoy writing (and coaching students to write) papers defensively because we fear the PC will have a "reviewer #2." And I don't love racing to hit a submission deadline (with the inevitable compromises that entails), waiting three months for a reply, learning that our paper didn't find its way to reviewers who understood it or shared its worldview, and then trying to figure out what minimal changes we can get away with before resubmitting to the next venue, as students become slowly demoralized by all the negative and delayed feedback. Being able to submit when ready, receive rapid engaged feedback from a knowledgable reviewer, clarify misunderstandings, and converge on something could be great.
- As a reviewer, I don't enjoy having a big pile of papers to review at the same time, especially when this takes me away from my real job of tending to a research group. I wish the intellectual effort that goes into understanding somebody else's work and writing a useful review could be part of the

mainline scientific discourse of our community; in the ARP system, if my students and I think hard about somebody else's work and put effort into writing a helpful and learned reply, that reply will add to the scientific discourse and we're accountable for it—some other research group is free to read it and tell us how we got it all wrong. In today's world, if a paper is getting rejected, I feel like I owe the authors a long review that covers the whole paper (because I only get one chance); it might be nice sometimes to give a shorter reply quickly that they can iterate on. And, from my time on the SIG Board and PCs, I've gotten depressed about all the LLM-aided misbehavior; e.g. the ACM has no functional process to deal with "content falsification" in an unpublished conference submission, so these situations can consume infinity hours of volunteer time. I miss the pre-COVID days when the PC was 40 people in a room for two days; I don't think the quality of deliberations is as good when the PC is 100+ people coming in and out of a Zoom call, and increasingly review rounds are passing by (or even the PC meeting itself) with reviews missing.

- As a conference attendee (and **early**-midcareer researcher!), I don't love finding myself among the most senior people in attendance. I think the post-COVID drop in attendance from senior-ish members of the community hasn't been helpful to younger folks and the community's development. And, while SIGCOMM certainly has some great talks, many talks at SIGCOMM do not seem that edifying to the attendee. It would be great to find a way to incentivize attendance from distinguished members of the community and to include more high-quality, thought-provoking talks in the program.

Conclusion: SIGCOMM ARP would be a "SIGCOMM beta" experiment, starting at SIGCOMM 2028, that lets authors submit year-round to the back of a particular reviewer's queue with a guaranteed latency bound for a reply. The authors would be anonymous, but their anonymized submission and the reviewer's identity and reviews would be public. The intention is to improve the author experience (by giving quicker replies from knowledgable reviewers), the reviewer experience (focusing on one paper at a time with feedback meant to help the authors), and the conference attendee's experience (improving attendance from senior members of the community and adding edifying talks to the program).

Acknowledgments: I'm grateful for conversations with the SIGCOMM Board and TSC, Justine Sherry, Anirudh Sivaraman, Scott Shenker, Aurojit Panda, Hari Balakrishnan, Marco Canini, Casey Henderson-Ross, Brad Karp, Philip Levis, and many others. They are not responsible for this proposal.

Related work (several of which I hope spawn their own "SIGCOMM beta" proposals):

- The SIGCOMM community's twenty-six "proposals advocating change" (2023). (See also the 2024 community session slides 44–53, summarizing progress on these.)
- Brighten Godfrey's "Iterative Reviewing for CS Systems & Networking Conferences" (2025)
- Fabian Bustamante, Nate Foster, Aurojit Panda, and Scott Shenker, "Accept Papers If And Only If There Is An Advocate" (2023).
- Nick Feamster, "Open Reviewing" (2023) & (2025)
- Russo & Storey, "From Passive to Participatory: How Liberating Structures Can Revolutionize Our Conferences" (CACM 2026)
- New Ideas in Networked Systems (2026–)
- Marco Canini, "Scientists Should Stop Writing Papers for Each Other" (2026)
- Proposals for CCR to serve as a continuous reviewing pathway for SIGCOMM.
- Journal of Systems Research (2021–)