More Streetlights, Please: On the Streetlight Effect in Networking Research

Vyas Sekar
Carnegie Mellon University
vsekar@andrew.cmu.edu

Abstract

A drunk man searches for his lost keys under a streetlight—not because that is where he dropped them, but because that is where the light is. Networking research, we argue, suffers from the same bias. As a community, we have systematically concentrated effort on problems that are well lit—problems formulated by industry or by eminent researchers that have well-defined measures of success and improvement potential—while leaving important unexplored territory in the dark. This perspective describes the *streetlight effect* through illustrative examples in networking research. We use a data-driven approach over historical publication data to demonstrate these streetlight effects on "hot" topics that have come (and gone) in the past. We qualitatively trace the persistence of this phenomena to the incentive structures surrounding authors, reviewers, program chairs, and funding agencies. We suggest, perhaps counterintuitively, that the solution is not to avoid streetlights, but to *set up more diverse streetlights*. We discuss some proposals to implement this idea, including several well-documented ideas to increase acceptances, explicitly encourage discussions with other communities, and explicitly celebrating research efforts that enable new areas to thrive. Admittedly, at the end of this perspective, we may have more questions than answers; but avoiding these questions is arguably more dangerous than having incomplete answers.

CCS Concepts

• **Social and professional topics** → *Socio-technical systems*;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
HotNets '26, Salt Lake City, UT, USA
© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-3029-0/2026/11
https://doi.org/10.1145/3847290.3859956

ACM Reference Format:
Vyas Sekar. 2026. More Streetlights, Please: On the Streetlight Effect in Networking Research. In *The 25th ACM Workshop on Hot Topics in Networks (HotNets '26), November 16-17, 2026, Salt Lake City, UT, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3847290.3859956

1 Introduction

There is an old joke about a drunk man who is searching for his keys under a streetlight. A passerby offers to help, and after several minutes of fruitless searching asks: "Are you sure you dropped them here?" The drunk replies: "No—I dropped them over there in the dark. But the light is so much better here."

This joke is not about drunkenness. In the philosophy of science it is known as the *streetlight effect*: the systematic bias toward research problems that have been previously well formulated and admit tractable measurement, regardless of whether those problems are important or if the effort invested is worthwhile. In other words, it is about the rational tendency to concentrate effort where effort and reward are tangible—not necessarily where it is most needed.

Science under the streetlight is not bad science. The techniques, artifacts, and advancements are real. But science that *only* happens under the streetlight gradually loses the ability to ask hard questions and risks developing systemic blind spots. The attendant research infrastructure, while an enabler of rigorous science, also perversely becomes a "gravitational black hole" that distracts people.

The history of networking research is littered with streetlights that burn bright for a while, attracting the attention of the best minds in the community. For instance, topics like ATM, multicast, P2P, and active networks burned bright for a brief while before being jettisoned by the community. More recently, we see new hot topics at the intersection of networking-for-AI and AI-for-networking. Some other topics like congestion control appear to be perennially popular streetlights. These problems are set up by either by industry efforts pointing to a real world problem or by popular papers from pre-eminent researchers. This in turn inspires the community to build datasets, simulators, testbeds, analytical models, workload generators, and benchmarks providing a rich substrate for enabling further research in this specific

area. These are genuine achievements, but could cause more harm than good in the long term.

This perspective is far from novel (e.g., [8]). In their call to action, the 2026 Hotnets chairs [5] explicitly ask whether "too much of our work is drifting toward short-term industry questions," and what the community should be teaching the next generation. In an earlier call for change, Shenker [9] identifies a related pathology; i.e., the community risks becoming unable to imagine what comes next—having nurtured the modern Internet, and finding itself trapped by it. The goal of writing this paper is to rekindle the discussion at Hotnets and beyond. This is also a critical and imminent time to begin this conversation in earnest as AI-driven acceleration of research and publication makes taste and judgment in problem selection and formulation ever more important than in the past.

In §2 we use a data-driven approach to illustrate the streetlight effect through some representative examples in networking research. In §3 we model the incentive landscape—authors, reviewers, program chairs, and funders—to understand why the effect persists even when individual researchers recognize it. In §4 we discuss ideas for a constructive response that do not rely on individual courage, but a community investment in planting more streetlights.

2 A Data-Driven Perspective On The Streetlight Effect

In this section, we present a data-driven analysis on SIGCOMM proceedings over the last three decades to illustrate the streetlight effect.

2.1 Methodology and Limitations

We built a simple "vibe coded" tool to automatically scrape ACM SIGCOMM main conference metadata from DBLP. Given the paper titles, we use a human-in-the-loop annotation approach that clusters the titles into research topics, refines classifications with an LLM agent, and produces an interactive timeline visualization. The code, data, and agent rules are all publicly available [1]. We acknowledge this methodology can be improved on multiple fronts; e.g., considering more SIGCOMM/Hotnets related conferences, improving the paper classification (e.g., including paper contents as part of the topic classification), and considering submissions not just acceptances.

2.2 Illustrative Examples

Figure 1 shows a bird's eye view of the temporal hotness of different streetlights between 1990-2025. We use this visualization and a few select examples to discuss some of the streetlight effects we see over time.

- *Congestion Control:* TCP congestion control is perhaps one of a few streetlights that do not seem to burn out. (Another such area is measurement and troubleshooting discussed below.) The metrics are defined, there are mature code bases, and new problems seem to arise as new workloads emerge. It peaked in the 2000s, saw a brief downturn, and has seen a mild resurgence in the datacenter era. That said, it does raise two questions: 1. Are we hitting diminishing returns investing in this area? 2. Do we have disproportionately many researchers working in this space?
- *Measurement and Troubleshooting:* The other perennial streetlight seems to be measurement and troubleshooting. This is arguably a streetlight that will and should continue to persist. As new workloads and usage patterns emerge, understanding workloads and failure patterns is key to improving systems. In some sense, this is an "umbrella" category since the specific topics that are being measured do evolve with time. For instance, topics like packet arrivals and long-range dependence seemed popular in the 90s but newer studies focus on cloud- and Internet-scale applications.
- *RDMA:* Even though the idea was posited as early as 1995 [2], RDMA was a bit of an outsider to the SIGCOMM community until 2014. This is perhaps a negative example of the streetlight effect where RDMA was not a popular topic in the early 2000s in the SIGCOMM community despite its growing adoption in HPC deployments. Networking research, however, did rectify that omission starting 2014. Under this bright light, the community has produced sophisticated work on congestion control for lossless fabrics, flow scheduling, and NIC offload.
- *Software-Defined Networking:* SDN arrived as the culmination of several years of evolution to decouple the control plane, centralize intelligence, make networks programmable [7]. The intellectual provenance was compelling and the attendant open-source contributions made it easy for researchers to evaluate research ideas in a domain that may have previously been restricted to vendors. While we saw a peak in SDN contributions in research around 2015-2016, it has since become mainstream and does not appear as an independent research agenda. That said, ideas and concepts from SDN are pervasively found in the research and deployment.
- *Programmable Switches:* Starting from the Forwarding Metamorphosis paper [3], the area of programmable switches saw a flurry of activity post 2014. Again, this was enabled by an ecosystem of open source tools and support from hardware vendors to enable research in the area. The activity in the area has died down

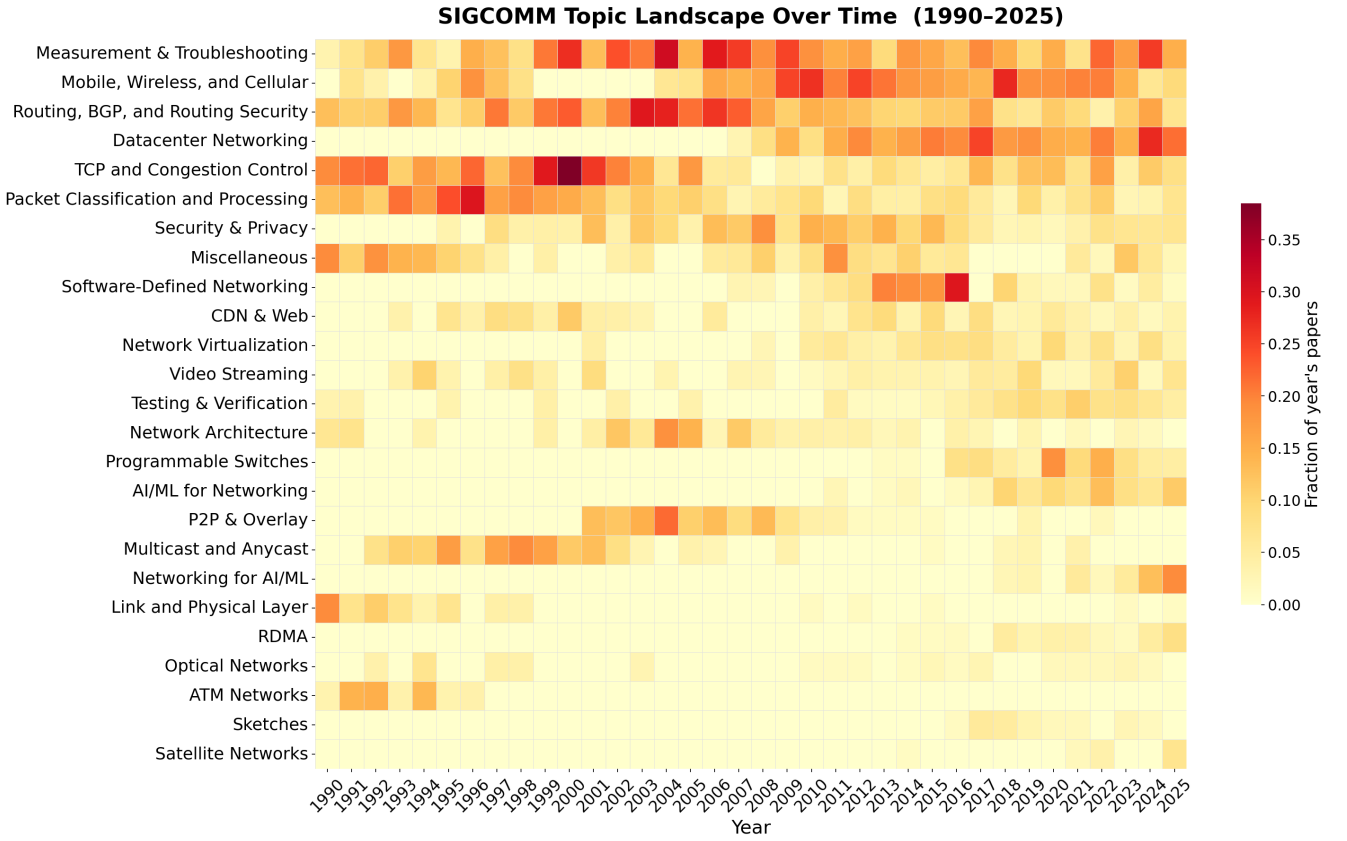


Figure 1: A historical view of the topics that have appeared in SIGCOMM from 1990-2025. We see some clear patterns where topics burn brightly for a bit and then die out. While seen over the entire history we have a diversity of streetlights at any given point a few topics seem to dominate the community attention.

somewhat, perhaps coincident with Intel "sunsetting" the effort.

- *Networking for AI:* The rapid growth of AI training and inference workloads has produced a new streetlight: scheduling for GPU clusters, network optimization for all-reduce collectives, and KV-cache management for inference. The proliferation and effort invested in AI workloads and benchmarks in adjacent communities makes it further easier to evaluate quick wins in this space. Unsurprisingly, we seem to be hitting a peak in terms of number of accepted papers on this topic.
- *AI for Networking:* A natural dual of the networking for AI is the application of machine learning to networking. While this idea is far from new, the advent of deep learning, deep RL, and new LLM/agentic capabilities is creating a new "cottage industry" in this space. That being said, compared to other applications of ML, we do not yet have standard datasets, benchmarks, and evaluation frameworks. This will likely take a few iterations to converge on mature domain-specific benchmarks for various tasks.
- Streetlights that have burned out: Along the way, we do see topics that seemed to have lit up briefly and then died out; e.g., ATM, Multicast, P2P, Link and physical layer, and active networks. While they were around, they did see a vibrant active effort and discussion, even if some of the human capital spent in these directions were futile in hindsight. Nevertheless, even the futile-in-hindsight streetlights do have value in the grander scheme of things as they serve as a valuable training ground for future systems and networking researchers.

2.3 An alternative viewpoint

Lies, damned lies, and statistics.

In the interest of intellectual honesty and completeness, it is only right that we also acknowledge an alternative viewpoint that has come up in discussions around hot topics and herding effects at various community discussions. A careful

reader looking at this same data can also potentially come to a positive conclusion— historical data confirms that the networking community is vibrant, evolving with time, picking relevant topics, and jettisoning obvious dead ends.

This is a valid counterpoint reading the same data. That being said, we discuss a few potential counterarguments to this alternative viewpoint. At the same time, however, we do see examples where we are "late to the game" (e.g., RDMA) or focusing excessive attention on hot topics (e.g., Networking for AI), at the cost of others. Furthermore, any data-driven approach has the inevitable bias of things that are not in the data. there is the obvious question of topics that are not even showing up on this list; e.g., technologies for resource-constrained settings or niche deployments (e.g., trading, disaster relief, home networks). Perhaps one could revisit this analysis with respect to *submissions* not *acceptances* to reveal even stronger biases. Note that even that could be an understatement as authors may themselves filter out topics if they think SIGCOMM would not think it is cool or trendy.

For the rest of the paper, we will conservatively assume that we do indeed see the streetlight effect, try to discuss why it occurs, and how to avoid it.

3 Ecosystem Issues That Perpetuate the Streetlight Effect

In some sense, the streetlight effect is an emergent property of an ecosystem (Figure 2) in which rational stakeholders—authors, reviewers, program chairs, and funding agencies—each responding sensibly to their mandates, can collectively produce a suboptimal outcome.

3.1 Authors: the credentialing trap

Researchers submit papers for a mixture of motives: *fame* (reputation, citation counts, community standing), *edification* (education the community on new and rigororus approaches), *joy* (the intellectual satisfaction of solving a hard problem), and *credentialing* (satisfying the requirements of a PhD program, a tenure case, or a grant renewal). Some or all of these push toward the streetlight.

Fame accrues to work that is palatable and legible to a wide audience: a paper that improves throughput on a standard benchmark by 30% is immediately comprehensible to any networking researcher. Edification is easiest when a problem is well-defined enough to know when it has been solved. Credentialing requires a paper, a paper requires an evaluation, and an evaluation requires well-defined criteria—which requires a streetlight.

The incentive to explore unlit spaces or to plant new streetlights is fragile. The publish-or-perish atmosphere leads researchers to alter their research to increase the chances of their work being accepted, pushing toward papers with narrower scopes and more quantifiable outcomes [9]. This is the credentialing trap: the career incentives meant to reward great science end up rewarding *easily measurable* science, which is not the same thing.

3.2 Reviewers: the legibility bias

Reviewers are expected to act as the community's gatekeepers or quality-control mechanisms. A reviewer who receives a paper with a clean experimental evaluation, standard baselines, and well-understood metrics can assess it quickly and with confidence. A paper that proposes a new problem, defines a new metric, and argues from first principles without a full evaluation requires a different kind of engagement—one that takes more time and carries more uncertainty about the correctness of the reviewer's judgment.

The rational response for a busy reviewer is to apply higher scrutiny to work that is harder to evaluate, and to reward work that speaks the community's established evaluation language. The aggregate effect is a systematic bias against work that ventures outside the existing measurement framework.

Shenker's diagnosis captures the core problem: conference acceptance is a *distorted* signal because "program committees will always have a hard time rejecting well-executed papers with limited scope" [9]. Papers that *define* new streetlights—new datasets, new benchmarks, new metrics—are especially vulnerable as they cannot yet demonstrate results and look underspecified by the (perhaps flawed) peer review standards biased toward existing problems and infrastructures.

3.3 Program Chairs: the false-negative problem

Program chairs face a different but related set of pressures. They are responsible for a program that will be judged by the community—which means a program organized around the community's current interests, evaluated by the community's existing criteria, rather than long term outcomes that move the community's scientific mission forward. That is, it systematically privileges the established streetlights.

Others have pointed out this core asymmetry directly: the risk of *false negatives*—incorrectly rejecting papers that are genuinely valuable—is more detrimental than the risk of false positives [5]. Interventions such as "sufficient support" (e.g., a single strong champion or two positive champions), while a good step forward, may not fully address the legibility problem. A paper exploring unmeasured territory is still evaluated against implicit norms developed for a different kind of contribution.

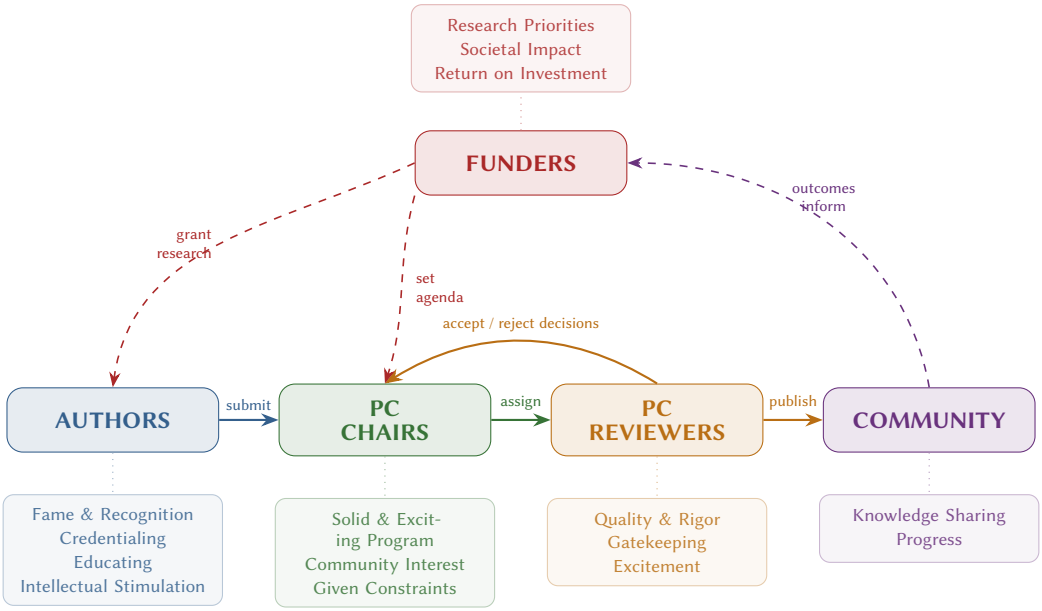


Figure 2: High level view of the various stakeholders and their motivations that end up perpetuating the streetlight effect

3.4 Funding Agencies

One natural cause of the streetlight effect could be funding incentives. Funding agencies may choose to invest more resources in hot topics (e.g., AI, quantum, Rust at the time of writing this paper in 2026) that create a herding effect. (That said, we should acknowledge that innovative technologies like the Internet or self-driving cars may have never manifested if not for a concerted and concentrated funding effort in a specific biased area.) Alternatively, the funding reviewing process may have similar biases to conference reviewing and may prioritize specific well-defined topics where intellectual contributions are easy to define (e.g., small proposals on congestion control are easier to justify as novel and safe to fund).

3.5 The Self-Reinforcing Cycle

Together, these pressures form a self-reinforcing cycle. Junior members of the community are forced into a (potentially degenerate) "RL" loop learning what gets accepted by looking at the current program and naturally try to produce more of the same. Reviewers develop expertise in evaluating what gets submitted; they develop less tolerance for departures from the tried-and-tested formulas. Program chairs optimize for the community's existing interests and those interests are shaped by what enters the submission pool to be reviewed. For each of the already lit streetlights, the attendant benchmarks, codebases, and datasets grow more sophisticated with each iteration, making the already lit region ever more attractive relative to unexplored spaces. Conference acceptance thus becomes "noisy, distorted, and distorting" [9]. Researchers optimize for the signal, which means optimizing for the benchmarks, which means working under the existing streetlights.

4 Incentivizing More Diverse Streetlights

To avoid this streetlight effect, one can hope that individual researchers have the courage to buck popular trends and work on hard, unlit problems with questionable reward. Unfortunately, hope or courage are neither good nor scalable strategies. What we need, instead, is a systemic change. We propose, perhaps counterintuitively, that to avoid the streetlight effect we need to *build more streetlights*. The question then is how we can systematically do this. Referring back to Figure 2, we need to think about the kinds of interventions that are possible at each of these stakeholders. In what follows, we discuss a few complementary (not necessarily novel) approaches that could help.

- **Plenary sessions from external experts:** The networking research community may have become unusually insular for a field whose subject, namely the Internet, is a consequential and societally relevant artifact. The problems that practitioners, healthcare workers, educators, city planners, climate scientists, and regulators face when interacting with networked systems are largely absent from the program. We propose dedicating a portion of each conference program to invited

talks from practitioners and domain experts *outside* the networking research community, tasked specifically with describing problems they face that current networking research cannot address. The speaker need not have a technical solution—the most valuable contribution is an articulation of the problem. This is a mechanism for identifying *new streetlights*; i.e., places where a modest investment could unlock a significant new research agenda.

- **Celebrate new research enablers:** The most durable contribution a researcher can make is often not a new system or algorithm: it is often creating something that enables others to build on. We do have examples of this both in networking (e.g., PlanetLab [4], Emulab [10]) and beyond (e.g., ImageNet [6]). Yet our current structure of reviewing provides little reward for this kind of hard infrastructure work. Networking research should explicitly solicit and elevate contributions that define new infrastructure, measurement frameworks, metrics, and datasets even when no results yet exist under those frameworks. A paper that argues convincingly that a new metric is needed, defines it precisely, describes how it could be measured, and identifies the research questions it would enable is doing important community work.
- **Funding diversity:** While it may be infeasible to decouple research from funding, one thing we could do is diversify fundamental research funding; e.g., philanthropic foundations with different mandates, internal funding from universities for topics that buck the norm, or enabling funding program managers to have some discretion.
- **Change review criteria on correct vs. novel:** As others have also observed, today the reviewing process treats a paper that is poorly executed and a paper that is exploring unmeasured territory somewhat similarly. A paper exploring unmeasured territory should be assessed on "the originality and substance of the perspective, the strength of the argument, and the potential to influence future research" [5].
- **Diversify reviewer pool to include non-networking experts:** The review ecosystem's bias toward established streetlights is amplified by homogeneity. A program committee composed entirely of networking researchers trained classically will, rationally and without bad intent, evaluate contributions against established norms. Work that imports problems or ideas from adjacent fields will be opaque to reviewers fluent only in the home domain. First, deliberately include reviewers from adjacent communities—distributed systems, ML systems, security, economics, measurement science—as regular program committee members, not occasional experts. Reviewer diversity does not guarantee better decisions, but it can expand the set of problems that are *legible* to the review process.

5 Discussion

We acknowledge a few other issues that have been raised in discussion with different members of our community.

Case for fun/interesting streetlights. It is possible that as a community we are taking ourselves a bit too seriously. Maybe to get new streetlights, we should accept papers that are either fun or enjoyable because they make people think about new things, rather than only those that are technically perfect implementations of the normal form.

Fostering subcommunities. An alternative is to encourage subcommunities within networking to exist and thrive. For instance, we have a thriving IMC community bringing together researchers focused on aspects of Internet measurement. Do we need similar subcommunities to foster more streetlights and have a "best of X' session at the main community conference?

What makes a new good streetlight. As a reviewer, is it possible to distinguish an important streetlight from a weak paper avoiding hard evaluation? Similarly, how do we assess if we are creating relevant streetlights in the long term?

Incentives for junior researchers. Junior researchers (graduate students, postdocs, early career assistant professors) are arguably most likely to uncover new and ambitious streetlights unencumbered by the baggage of history. But somewhat perversely, the system is engineered to prevent them from taking the risk in investing time and effort on new streetlights and may in fact negatively reward the risk takers (e.g., paper and grant rejections, tenure denials).

Retrospectives on why streetlights pop up and die. It may be useful for the community to explicitly do retrospectives on why some streetlights popped up (e.g., funders, societal change, new tech enablers) and why some failed (e.g.., funding stopped, problem was solved, problem was irrelevant) and why others have been more permanent.

Learning from other communities. The streetlight effect or proposed solutions are definitely not unique to networking research, though there may be some aspects (e.g., scale, real users) that are specific to networking. This raises the question of whether we can learn from other communities (e.g., physics, economics, mathematics, biology, other areas of computing/engineering) that may have avoided the streetlight effect. Or alternatively, are there lessons we learn from our experiments that we can share with these communities?

Are streetlights really bad? One counterargument is that well-defined streetlights help train the next generation of researchers in the scientific method and best practices. One could argue that this "human capital"; i.e., future researchers are the actual product of the endeavor whose mission is to find and plant their own streetlights, rather than the specific research papers or artifacts under the streetlights. Seen in this light, maybe the streetlight effect is not so bad after all!

6 Conclusions

In some sense, none of the problems or ideas presented here are new. The goal of this perspective is to rekindle the broader debate about our scientific mission.

Networking research has produced remarkable high-impact results under its existing incentive structure. We are not arguing that this work should stop. This perspective is arguing that the community should invest, with more seriousness in ecosystem support that will let us illuminate new territory.

The streetlight effect is not a law of nature. It is a consequence of the incentive structures, evaluation norms, and measurement infrastructure we have built. We can build different ones. We should.

Acknowledgments

This paper became stronger with feedback from a number of colleagues. The author thanks Aurojit Panda, Justine Sherry, and Keith Winstein for their constructive feedback on an earlier version of the draft. The author thanks Milind Srivastava, Lesley Zhou, and Marko Morrison for their feedback representing the (underrepresented) voice of graduate students that is often missing in broader debates around the types of change we would like to see in the conference, reviewing, and publication processes. The author also thanks many colleagues for interesting hallway discussions and ideas that helped shape this paper over many years. (As they say, good artists copy, and great artists steal and hide their sources. In this case, the author is neither an artist nor good nor great, merely forgetful.) The author thanks the anonymous Hotnets reviewers and the shepherd Venkat Arun for taking an unconventional chance on a paper that admittedly does not say anything surprising, but one that could spark an interesting discussion.

References

[1] [n. d.]. Data and code for streetlight effects analysis . https://github.com/vyassekar/hotnets2026_streetlight/.

[2] Ken Birman. 2016. A short history of remote DMA networking. https://thinkingaboutdistributedsystems.blogspot.com/2016/12/rdma-1-short-history-of-remote-dma.html. Accessed: 2026-07-07.

[3] Pat Bosshart, Glen Gibb, Hun-Seok Kim, George Varghese, Nick McKeown, Martin Izzard, Fernando Mujica, and Mark Horowitz. 2013. Forwarding metamorphosis: fast programmable match-action processing in hardware for SDN. *SIGCOMM Comput. Commun. Rev.* 43, 4 (Aug. 2013), 99–110. doi:10.1145/2534169.2486011

[4] Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. 2003. PlanetLab: an overlay testbed for broad-coverage services. *SIGCOMM Comput. Commun. Rev.* 33, 3 (July 2003), 3–12. doi:10.1145/956993.956995

[5] P. Costa and M. Schapira. 2026. A backward-compatible protocol upgrade for HotNets. *arXiv preprint arXiv:2606.25786* (June 2026).

[6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255. doi:10.1109/CVPR.2009.5206848

[7] Nick Feamster, Jennifer Rexford, and Ellen Zegura. 2014. The road to SDN: an intellectual history of programmable networks. *SIGCOMM Comput. Commun. Rev.* 44, 2 (April 2014), 87–98. doi:10.1145/2602204.2602219

[8] J. Rexford and S. Shenker. 2022. Answering three questions about networking research. *ACM SIGCOMM Computer Communication Review* 52, 1 (2022), 42–44.

[9] S. Shenker. 2022. Rethinking SIGCOMM's conferences: Making form follow function. *ACM SIGCOMM Computer Communication Review* 52, 4 (2022), 26–30.

[10] Brian White, Jay Lepreau, Leigh Stoller, Robert Ricci, Shashi Guruprasad, Mac Newbold, Mike Hibler, Chad Barb, and Abhijeet Joglekar. 2003. An integrated experimental environment for distributed systems and networks. *SIGOPS Oper. Syst. Rev.* 36, SI (Dec. 2003), 255–270. doi:10.1145/844128.844152