

ACM SIGCOMM Student Mentoring Column

Various Student Questions

Dear students: I hope you have been enjoying reading the column. This column is similar to the previous one: it attempts to address a few more of your questions. Thanks for the great questions; keep them coming! Again, many thanks to Brighten Godfrey (UIUC) and Vyas Sekar (CMU) for contributing their thoughts.

On Research Topics in Networking, and Reading Papers

Q: *“What are the top 10 papers any networking student should know about?”*

A: This is a difficult single question to answer, especially as the networking discipline has become very diverse. It really depends a lot on the “flavor” of networking you may be interested in -- systems vs theory, wired vs wireless, applications vs. lower level, etc. No matter the flavor, here are some ways of finding those top papers to read:

- One starting point is to look at what some eminent networking researchers have come up with in the past. There are many such example lists, including those compiled by Jim Kurose [1], Jon Crowcroft [2], Mostafa Ammar [3], George Varghese [4], and Jen Rexford (with a focus on networking theory papers) [5].
- Another idea is to look at papers that have received major awards. An example is the set of “test of time” award papers at SIGCOMM and NSDI. These are papers that were well ahead of the pack when they were published, but have had significant real world impact over the years. Similarly, it is worth looking at papers that have won IRTF’s Applied Networking Research Prize. These reflect research ideas with significant potential for impact on practice.
- A final idea is to look at the Graduate Networking Classes offered and the recommended readings therein, especially on “classical” topics. You are welcome to check out UW-Madison’s graduate networking [6].

Q: *“How can I get more efficient at reading papers?”*

A: Keshav’s “How To Read a Paper” [7] presents a great prescriptive “algorithm” for becoming more efficient at reading a paper. I won’t spoil the surprise for you, but the gist is to do three passes in progressively increasing depth of understanding. The depth to which you need to understand the details of the paper depends on the context/reason you are reading it. For instance, if you are reviewing for a conference/journal you really need to do an in-depth read during the third pass. The same applies if you are planning to work in the same general area as the paper, or if the paper is highly relevant to your current work. In both cases, it may also help to take detailed notes and jot down your thoughts/critiques during the in-depth pass. This may seem onerous at the beginning, but over time you will notice that you’re getting very efficient with each pass taking lesser time. On the other hand, if you are doing a quick survey to understand the state-of-art then you may do only the first/second passes. Generally speaking, the only way to become more efficient at something is practice. So read as many papers as you possibly can.

A Few More Publication-related Questions

This appears to be a popular topic, with many different questions. I’ve attempted to address a handful of them below, and will return to the others in future columns.

Q: *“Is X a good venue to publish at?”*

A: The student had actually asked about a specific “X”, which I decided to anonymize: the purpose of this column is not to answer questions about particular venues, but to offer general advice. To that point, here is a general set of guidelines in choosing what conference venues to publish in:

- Are the leading researchers in your field publishing at this venue? Look at prior years’ programs to understand this.
- How visible/impactful are papers at this venue? Are these papers cited a lot? Lookup papers published at the venue over the past 2-3 years in Google Scholar. Were some of the “seminal” papers in your field published here? Are these papers you’d put on your must read list?
- Do your peers/advisor publish or want to publish here? It doesn’t hurt to ask directly.



• Is the PC composed of well-recognized members of the community (e.g., people whose papers you read)? In addition to hinting at the quality of reviews and feedback you would get, this also tells you whether the venue will be well attended and by relevant folks (because in most cases members of the PC do submit to, and attend, the conference).

Q: *“Why do networking conferences seem to go through these “fads” every few years – multicast, DHTs, peer-to-peer, datacenter networking, and now SDN?”*

A: Hey, don’t leave out quality of service, ad hoc networking, and the Internet of Things! But seriously, this happens in most fields of research and technology, not just networking. Hence the “hype cycle” [8]. In general, these things are also driven by where the industry is headed or most focused, technology trends, as well as funding agencies’ areas of priority. If you want to know what these “hot topics” are at the moment, there are a few different sources you can peruse: (a) You can pick up almost any networking trade magazines and you will find the relevant buzzwords there! (b) Read the most recent call for proposal from the US’s National Science Foundation, or similar organizations at other countries; (c) read conference calls for papers (CFPs).

Q: *What are some good strategies to ensure that my work/publication has “real world” impact?*

A: This is an excellent, and a very important question. Work that has demonstrable real world impact is much more valued, e.g., in job search, than work that did not go beyond a paper. There are a few ways that, in conjunction, can help your work achieve impact:

- Fully build out your prototype systems, and release software for public use. This will ensure other people, including those in the industry, build on your work. In cases where folks from the industry read your paper and want to license your work, they will often want a live demo.

Having the system built out will help you be prepared for that.

- Invite yourself to give industry talks. This is a great way to publicize your work to the relevant groups in key companies, as well as to obtain pointed, relevant and practically important feedback.

- Sign up to speak at, or conduct demonstrations at, operator fora, such as NANOG, ONUG and ONS. In addition to publicity and feedback, you may get real world operational insight – and datasets! – that open your eyes to important extensions you can pursue for your system.

References:

- [1] “10 networking papers: recommended reading”, Jim Kurose. ACM SIGCOMM Computer Communication Review, Volume 36 Issue 1, January 2006, Pages 51-52.
- [2] “10 networking papers: recommended reading”, Jon Crowcroft. ACM SIGCOMM Computer Communication Review, Volume 36 Issue 2, April 2006, Pages 31-32.
- [3] “10 networking papers: a blast from the past”, Mostafa Ammar. ACM SIGCOMM Computer Communication Review, Volume 37 Issue 1, January 2007, Pages 57-59.
- [4] “10 network papers that changed the world”, George Varghese. ACM SIGCOMM Computer Communication Review, Volume 37 Issue 5, October 2007, Pages 51-53.
- [5] “My Ten Favorite “Practical Theory” Papers”, Jennifer Rexford. ACM SIGCOMM Computer Communication Review, Volume 38 Issue 2, January 2008, Pages 41-43.
- [6] UW-Madison Graduate Computer Networking, <http://pages.cs.wisc.edu/~akella/CS740/S15/>
- [7] “How to Read a Paper”, S. Keshav. ACM SIGCOMM Computer Communication Review, Volume 37 Issue 3, July 2007, Pages 31-33.
- [8] Hype Cycle. https://en.wikipedia.org/wiki/Hype_cycle

Aditya Akella

University of Wisconsin-Madison, USA

