

10 Papers for the Ph.D. Student in Networking

Craig Partridge
BBN Technologies

One of the challenges I remember from my days as a Ph.D. student was the tremendous struggle to figure out what papers I should read: what key ideas did I need to understand to feel I was even vaguely qualified to do work in networking?

The idea in this essay is to try to answer those questions for the Ph.D. student of today while fitting within the limit of ten papers. Two of these works, alas, are doctoral dissertations, not papers but I am hoping the *CCR* editor will let that slide...

1. J.H. Saltzer, D.P. Reed and D.D. Clark, "End-to-End Arguments in System Design," *ACM Trans. on Computer Systems*, Vol. 2, No. 4, Nov 1984, pp. 277-288

There are not many papers on communications architecture. This paper is the best known and it rewards multiple readings, largely because it is so hard to apply. It has spawned multiple interpretations (e.g., many argue that Isenberg's stupid network [1] is a consequence of the end-to-end argument) and even has a retrospective analysis from twenty years later [2].

2. A. Demers, S. Keshav and S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," *Proc. ACM SIGCOMM '89*, pp. 1-12.

I love this paper because it turns paradigms on their head. Wasn't the whole point of packets to get rid of fixed bandwidth channels and bit-by-bit multiplexing? In this wonderful paper, the authors show how to map between the bit-by-bit world and the packet-based world. In so doing, they pointed the way to offering the delay and bandwidth guarantees of a bit-by-bit environment in a packet-switched world [3].

3. S. Floyd and V. Jacobson, "The Synchronization of Periodic Routing Messages," *IEEE/ACM Trans. on Networking*, Vol. 2, No. 2 (1994), pp. 122-136.

You'll never think the same way about clocks and timers and packet interactions again. And your protocols will be the better for it.

4. V.G. Cerf and R.E. Kahn, "A Protocol for Packet Network Interconnection," *IEEE Trans. on Communications*, Vol. 22, No. 5, May 1974, pp. 637-648.

The key reason to read this paper is not its historical significance but for an understanding of how ideas evolve. This paper launched the Internet. Yet it has no Internet

Protocol (IP), is unsure about where to do fragmentation and reassembly, and has no congestion window. So I recommend reading this paper as the starting point of an exploration of how and why TCP evolved into TCP/IP in the 1970s. To watch TCP evolve during its formative years, follow reading this paper with references [4-9].

5. A. Demers, D. Greene, C. Houser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart and D. Terry, "Epidemic Algorithms for Replicated Database Maintenance," *Proc. ACM PODC '87*, pp. 1-12.

This is one of the later papers about the Grapevine distributed system [10], a system which deserves more attention than it gets. Of all the Grapevine papers, I find this paper the most satisfying because, beyond explaining some of Grapevine, it imparts a key message: namely that many problems can be solved if you're willing to permit a little bit of inaccuracy in the system. It is an important approach that has been used profitably in protocol verification [11] and packet tracing [12].

6. J. Keilson and L.D. Servi, "A distributional form of Little's law", *Operations Research Lett.* 7 (1988) 223-227.

This paper is a queueing theory paper. Surprisingly you can still read Kleinrock's queueing theory textbooks from the 1970s and learn most of what you need to know about the field. But there are a few subsequent innovations one should be aware of. Wolff's PASTA theorem is probably the most important [13].

Keilson and Servi's paper is a very practical result you are very likely to miss. It is a short cut for converting from the formulations of the inter-arrival times and service times to the distribution functions for the time in system and time in queue. The rule applies for a wide range of systems with Poisson arrivals.

7. M. Waldvogel, G. Varghese, J. Turner, and B. Plattner, "Scalable High Speed IP Routing Lookups," *Proc. ACM SIGCOMM '97*, pp. 25-36.

Devising new algorithms to solve vital problems is an important part of networking research. This is one of the prettier algorithm papers I know. It also made a difference. It was one of two papers in 1997 (the other was [14]) that gave the community the confidence that it could re-architect routers to accommodate data rates of tens to hundreds of gigabits per second [15].

8. M. Mathis and J. Mahdavi, "Forward Acknowledgement: Refining TCP Congestion Control," *Proc. ACM SIGCOMM '96*, pp. 281-291.

I'm always amazed at how much information you can glean from TCP segments and IP datagrams. This paper is a powerful example of just how much information a TCP sender has when it examines its state and the ACKs it is receiving. Other insightful papers of this ilk are [9] and [16].

9. Bruce J. Nelson, *Remote Procedure Call*, Ph.D. Thesis in Computer Science, Carnegie Mellon University, 1981.

My first advice to any Ph.D. candidate is to remember that the purpose of the thesis is to show you understand how to do original research in all its nitpicking detail. A corollary is that brilliance in a thesis is not required. That said, there's no harm in seeing what a brilliant Ph.D. thesis looks like. This is the thesis that solved most of the problems in implementing remote procedure call. Nowadays people cite the paper with Andrew Birrell written a few years later [17] — but most of the key ideas are already here. And the sailing theme is a fun touch.

10. Chapter 2 of Stuart Cheshire, *Consistent Overhead Byte Stuffing*, Ph.D. Thesis in Computer Science, Stanford University.

Writing a prior work section of a thesis, or a paper, is always difficult. You have to show both that you understand the prior work and know how to place your work in the context of the prior work. In Stuart Cheshire's case, nobody had done a prior survey of the literature. Indeed, the topic area was sufficiently arcane that almost no one had written any papers. Yet he delved into software, patents and whatever else he could find. Then he developed a taxonomy that allowed him to classify all work using four yes-no criteria and an efficiency metric. The result is a chapter that takes up half the dissertation, but that's because after this chapter, Cheshire's work (a wonderful technique for byte-stuffing) is so clearly a contribution, not much more needed to be said.

References

1. D. Isenberg, "The Dawn of the Stupid Network," *ACM Networker*, 2.1, February/March 1998, pp. 24-31.

2. M. Blumenthal, D. Clark, "Rethinking the Design of the Internet: the End-to-end Arguments vs. the Brave New World," *ACM Transactions on Internet Technology*, Vol. 1, No. 1 (August 2001) pp. 70-109.

3. A. Parekh, *A Generalized Processor Sharing Approach to Flow Control*, MIT Ph.D. Thesis 1992.

4. R.S. Tomlinson, "Selecting Sequence Numbers," *Proc. ACM SIGCOMM/SIGOPS Interprocess Communications Workshop*, March 1975, pp. 11-23.

5. D.D. Clark, "Window and Acknowledgement Strategy in TCP; RFC-813" *Internet Requests for Comments*, No. 813, July 1982.

6. J. Nagle, "Congestion Control in IP/TCP Internetworks; RFC-896", *Internet Requests for Comments*, No. 896, Jan 1984.

7. P. Karn and C. Partridge "Estimating Round-Trip Times in Reliable Transport Protocols," *Proc. ACM SIGCOMM '87*, pp. 2-7.

8. C.A. Kent and J. Mogul, "Fragmentation Considered Harmful," *Proc. ACM SIGCOMM '87*, pp. 390-401.

9. V. Jacobson, "Congestion Avoidance and Control," *Proc. ACM SIGCOMM '88*, pp. 314-329.

10. A.D. Birrell, R. Levin, M.D. Schroeder and R.M. Needham, "Grapevine: an Exercise in Distributed Computing," *CACM*, Vol 25, No. 4, Apr 1982, pp. 260-274.

11. G. Holzmann, *Design and Validation of Computer Protocols*, Prentice-Hall, 1991.

12. A.C. Snoeren, C. Partridge, L.A. Sanchez, C.E. Jones, F. Tchakountio, B. Schwartz, S.T. Kent and W.T. Strayer, "Single-packet IP Traceback," *IEEE/ACM Trans. on Networking*, Vol. 10, No. 6, (Dec 2002), pp. 721-734.

13. R.W. Wolff, "Poisson Arrivals See Time Averages," *Operations Research*, Vol. 30, No. 2 (1982), pp. 223-231.

14. M. Degermark, A. Bronik, S. Carlsson and S. Pink, "Small Forwarding Tables for Fast Routing Lookups," *Proc. ACM SIGCOMM '97*, pp. 3-14.

15. C. Partridge P. Carvey, E. Burgess, I. Castineyra, T. Clarke, L. Graham, M. Hathaway, P. Herman, A. King, S. Kohlami, T. Ma, J. Mcallen, T. Mendez, W. Milliken, R. Pettyjohn, J. Rokosz, J. Seeger, M. Sollins, S. Storch, B. Tober, G. Troxel, D. Waitzman and S. Winterble, "A 50-Gb/s IP Router," *IEEE/ACM Trans. on Networking*, Vol. 6, No. 3 (June 1998), pp. 237-248.

16. S. Floyd, "TCP and Explicit Congestion Notification," *ACM Computer Communication Review*, Vol. 24, No. 5, Oct 1995, pp. 10-23.

17. A.D. Birrell and B.J. Nelson, "Implementing Remote Procedure Calls," *ACM Trans. on Computer Systems*, Vol. 2, No. 1, 1984, pp. 39-59.

Doctoral dissertations are available from ProQuest (www.proquest.com).

Many thanks to Jim Kurose and Tim Strayer for their helpful comments on this paper.