

10 Papers on Network Models

Matthias Grossglauser
School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL)
1015 Lausanne, Switzerland

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General

General Terms

Performance, Scalability, Topology, Routing, Control, Inference

Keywords

Network Models

In the first draft of my list, more than half the papers I had chosen had already been mentioned in the three lists that preceded mine (in particular, [5, 6, 7, 15, 16, 12, 10]). Therefore, I decided to narrow the scope of my list around a theme, following David Wetherall's example. If the idea of the CCR recommended reading list continues (which I hope), I would assume that at least some of the future lists will likewise be focused on particular subtopics or on particular viewpoints. My list focuses on network models.

Networking is an engineering discipline, and many of the key advances in networking came out of efforts with an applied flavor, favoring systems, architecture, and experimentation over mathematical modeling and formal methods of analysis. Often, major new applications did not even emerge from academic projects, but were developed by creative young people outside the research community. This is in contrast to other disciplines, e.g., signal processing and communications, that live and breathe through their models.

I believe that one of the reasons for this is the sheer complexity of the task at hand: building a global, scalable, robust, secure, evolvable communication infrastructure, satisfying a broad panoply of needs and interests, resisting a wide range of threats, absorbing new technologies and supporting new applications, is much more likely to fail because of an inadequate architecture than because of inadequate underlying models. The former can doom the whole undertaking (and we all know the acronyms lying dead by the wayside), while the latter is often only a matter of efficiency, of throwing more resources at the problem. Furthermore, Moore's law and related exponential trends move the goalposts so quickly that clean, relevant modeling questions are often hard to find.

I do believe, however, that models have something to say.

First, models allow us to reason about scale before we are able or willing to build a large-scale artifact. Second, models can reduce the complexity of a question, by forcing us to try to filter out the salient features of a problem at hand. Third, new and practically relevant ways of approaching and solving a problem can come out of modeling efforts. Fourth, models are a useful didactic tool for the study of the systems we are building, in analogy to the natural sciences.

Here, then, are some papers on network models that inspired me, and that I believe made a contribution beyond theory, by providing fresh insights or by showing the way to new solutions to old problems. The usual disclaimer applies: the list is heavily biased by my own interests. Enjoy!

- L. Kleinrock and F. Kamoun, Hierarchical Routing for Large Networks, Performance Evaluation and Optimization, *Computer Networks*, 1(3), pp. 155-174, January 1977.

A very prescient paper, recognizing the scalability issues that arise in routing for very large networks. The paper shows the benefits of a hierarchical address space to make routing scalable. Many seeds of key ideas of Internet routing can be found in this paper. I am particularly impressed by the scale of the networks studied in this paper (up to millions of nodes), which certainly made little practical sense in the seventies. Perhaps a lesson to behold for networking research today.

- W. E. Leland, M. S. Taqqu, W. Willinger, D. V. Wilson, On the Self-Similar Nature of Ethernet Traffic, *IEEE/ACM Transactions on Networking*, 2(1), pp. 1-15, February 1994.

This paper was a watershed event in network and tele-traffic modeling. It argued very convincingly that network traffic has self-similar structure, which was a baffling observation at the time. It spawned a furious debate about causes and models, which, interestingly, paralleled similar discoveries in several other fields, such as hydrology [9] and finance [13].

- Y. Vardi, Network Tomography: Estimating Source-Destination Traffic Intensities from Link Data, *Journal of the American Statistical Association*, vol. 91, March 1996.

This is the first paper I am aware of that formulates the estimation of a traffic matrix (the amount of traffic between source-destination pairs) as a statistical

inversion problem. This has given rise to a rich area of research for both active and passive measurements (e.g., [3, 14, 4, 19]). These results are of real practical significance to large network operators, given the lack of precise and complete traffic measurements available in the Internet today.

- F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability, *Journal of the Operational Research Society*, 49(3), March 1998.

This paper shows that a network can regulate traffic flows to achieve a notion of fairness. Furthermore, this can be achieved in a fully distributed manner, where the network and each user optimize a utility function. The only coupling between the network and each user is through prices. This elegant model provides strong theoretical justification for TCP-like congestion control, and was one of the first to put congestion control and network-wide fairness issues in a rigorous control-theoretic framework.

- J. Kleinberg, The Small-World Phenomenon: An Algorithmic Perspective, in *Proc. 32nd ACM Symposium on Theory of Computing (STOC)*, Portland, Oregon, May 2000.

This paper makes a wider point about large complex networks, but I believe that this research also sheds light on actual routing problems, in particular for wireless networks. Kleinberg studies the problem of routing with purely local information in a lattice with random shortcuts, inspired by models for small-world networks. He shows that the empirically observed ability of people to forward messages through a social network to a destination [18] is in fact quite surprising, and requires that the shortcuts follow a very specific distance distribution.

- P. Gupta and P. R. Kumar, The Capacity of Wireless Networks, *IEEE Transactions on Information Theory*, 46(2), pp. 388-404, March, 2000.

This paper introduces a model for large wireless ad hoc networks, where nodes are both terminals generating traffic, and potential relays that forward traffic on behalf of other nodes. The authors study the throughput capacity, i.e., the throughput achievable concurrently for all sessions, for large networks. They show that this throughput goes to zero as $O(1/\sqrt{n})$, where n is the number of nodes in the network. This paper has led to animated discussions about the viability of large ad hoc and sensor networks, and motivated many related studies to overcome this capacity limitation.

- S.-Y. R. Li and R. W. Yeung and N. Cai, Linear Network Coding, *IEEE Transactions on Information Theory*, 49(2), p. 371-381, February 2003.

This paper introduces the notion of network coding. In network coding, intermediate nodes in a network (linearly) combine incoming messages to produce new outgoing messages. The potential benefits of network coding are efficiency gains in terms of network throughput and delay, and robustness in situations where messages

may be lost, or where a receiver does not have precise control over which messages it receives. Among the drawbacks is the need for internal nodes to buffer packets and to compute linear combinations. Therefore, the most promising applications of network coding so far appear to be wireless networks and peer-to-peer file distribution. For a survey, see [8].

- P. Francis, S. Jamin, Ch. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang, IDMaps: A Global Internet Host Distance Estimation Service, *IEEE/ACM Transactions on Networking*, 9(5), pp. 525 - 540, October 2001.

This paper is one of the first to propose the use of techniques from metric embedding theory in order to offer a distance estimation service between any pair of nodes in the Internet. What is particularly interesting about this line of work is that it directly motivated metric embeddings with slack, a fresh approach on a classical problem in theoretical computer science [11, 1]. This has been a very active area recently, and interestingly, some of the new results based on beacon-based embedding seem to be directly inspired by the applied work in the Internet community. Another interesting recent application of embeddings is for geographic routing in wireless networks [17].

- Timothy G. Griffin, F. Bruce Shepherd, and Gordon Wilfong, The Stable Paths Problem and Interdomain Routing, *IEEE/ACM Transactions on Networking*, 10(2), pp. 232-243, April 2002.

Inter-domain routing remains, I believe, the most difficult and critical problem in Internet operation. While BGP is the de-facto standard for inter-domain routing, the expression of routing policies is not in fact based on a clean underlying model, and the practice of BGP configuration today is essentially entirely ad hoc. This paper introduces the stable paths problem, which allows us to reason about the interaction of BGP policies of different service providers. It is a first step towards an inter-domain routing methodology that is based on logic rather than witchcraft.

- L. Li, D. Alderson, W. Willinger, and J. Doyle, A First-Principles Approach to Understanding the Internet's Router-Level Topology, in *Proc. ACM SIGCOMM 2004*, Portland, Oregon, Aug.30-Sept.2, 2004.

This paper, while making a point about the structure of the Internet, is of wider significance. It shows that the class of graphs with a given degree distribution is very large, and contains graphs with very different properties. Specifically, the authors show that the claim that high-degree nodes form the center of the network, and that these nodes are therefore of particular significance for the robustness of the network, is not true. This misunderstanding arose in part because properties of scale-free graphs generated through preferential attachment were falsely believed to hold more generally [2].

Acknowledgments

Many thanks to Jim Kurose for feedback on a draft.

References

- [1] I. Abraham, Y. Bartal, T-H. Chan, K. Dhamdhere, A. Gupta, J. Kleinberg, O. Neiman, and A. Slivkins. Metric embeddings with relaxed guarantees. In *46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [2] A.-L. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286:pp. 509–512, 1999.
- [3] R. Cáceres, N. G. Duffield, J. Horowitz, and D. Towsley. Multicast-Based Inference of Network-Internal Loss Characteristics. *IEEE Transactions on Information Theory*, 45:2462–2480, 1999.
- [4] J. Cao, D. Davis, S. Vander Wiel, and B. Yu. Time-Varying Network Tomography. *Journal of the American Statistical Association*, December 2000.
- [5] D. D. Clark. The Design Philosophy of the DARPA Internet Protocols. In *Proc. ACM SIGCOMM '88*, pages 106–114, Stanford, CA, August 1988.
- [6] Rene L. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.
- [7] Rene L. Cruz. A Calculus for Network Delay, Part II: Network Analysis. *IEEE Transactions on Information Theory*, 37(1):132–141, January 1991.
- [8] C. Fragouli, J-Y. Le Boudec, and J. Widmer. Network Coding: An Instant Primer. *ACM SIGCOMM Computer Communication Review*, 36(1), January 2006.
- [9] H. E. Hurst. Long-Term Storage Capacity of Reservoirs. *Proc. American Society of Civil Engineers*, 116:770–808, 1951.
- [10] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability. *Journal of the Operational Research Society*, 49(3), March 1998.
- [11] J. Kleinberg, A. Slivkins, and T. Wexler. Triangulation and Embedding Using Small Sets of Beacons. In *FOCS*, pages 444–453, 2004.
- [12] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE/ACM Trans. on Networking*, 2(1):1–15, February 1994.
- [13] B. Mandelbrot. The Variation of Certain Speculative Prices. *Journal of Business*, 36:394–419, 1963.
- [14] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic Matrix Estimation: Existing Techniques and New Directions. In *Proc. ACM SIGCOMM '02*, pages 161–174, 2002.
- [15] A. K. Parekh and R. G. Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks - The Single Node Case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.
- [16] A. K. Parekh and R. G. Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks - The Multiple Node Case. *IEEE/ACM Transactions on Networking*, 2(1):137–150, April 1994.
- [17] A. Rao, S. Ratnasami, C. Papadimitriou, S. Shenker, and I. Stoica. Geographic Routing Without Location Information. In *Proc. ACM MOBICOM 03*, San Diego, CA, September 2003.
- [18] J. Travers and S. Milgram. An Experimental Study of the Small World Problem. *Sociometry*, 32(4):pp. 425–443, December 1969.
- [19] Y. Zhang, M. Roughan, C. Lund, and D. Donoho. An Information-Theoretic Approach to Traffic Matrix Estimation. In *Proc. ACM SIGCOMM 03*, Karlsruhe, Germany, August 2003.