Reasoning about Networks

Janne Riihijärvi and Petri Mähönen

Department of Wireless Networks, RWTH Aachen University
Kackertstrasse 9, D-52072 Aachen, Germany
email: {jar, pma}@mobnets.rwth-aachen.de

Abstract: We discuss the role of geometry and topology in reasoning processes related to wireless and wired communication networks. We focus especially on future research directions on enabling autonomic reasoning taking place within the networks themselves. This work is a part of our ongoing effort to better understand future "cognitive communication systems" that are likely to employ various adaptive techniques, also borrowing from artificial intelligence research community to optimise their behaviour in an automated manner. As the field is very young, we mainly attempt to succinctly review and organize the state-of-the-art, and identify areas in which further research is urgently needed.

1 Introduction

Networks are expected to become more and more heterogeneous and complex in the near future. Introduction of substantial wireless network components, especially wireless meshes and sensor networks, coupled with mobility support will also take network dynamics on a new level. In these types of environments effective self-organization mechanisms are a must if any degree of optimality is hoped for in terms of quality of user experience. For further discussions on future of networking as is relevant for the following discussion, see [MRPS04] and references therein.

Traditionally network self-organization has been accomplished in rather rudimentary manner, essentially relying on basic address/host autoconfiguration and adaptive routing. Protocols used have been designed to carry out relatively simple optimisation tasks, based on classical graph models of networks. A good example is given by the various ad hoc routing protocols suggested. Instead of carefully considering energy efficiency, expected transmission times, and other performance metrics that are of acute interest to the users, simple hop-count minimization is typically performed.

In this paper we discuss potential approaches that can be adopted as a basis for future self-organization mechanisms. We stress the importance of developing realistic network models and *abstractions* of generic characteristics found in existing networks, and successful models. For large networks operating on high-level abstractions is in our view the only way to solve otherwise difficult scalability problems. We believe that "cognitive" techniques, borrowing from artificial intelligence community will play a central role in future self-organization mechanisms. For first suggestions along these lines, see [CPRW03].

2 Models, Abstractions and Network Dynamics

Put roughly, we can view the study of network features from two perspectives. On the one hand we can develop *models*, algorithms that generate, for example, suitably labelled graphs resembling connectivity graphs of existing networks from small number of input parameters. This approach has its main use in generation of network topologies for simulation purposes. On the other hand, more fundamental is, in our view, the study of *abstractions*, general characteristics that several existing networks and models thereof apparently share.

The abstraction with longest ancestry is undoubtedly that of a *small world network*. Originating from Milgrom's study of social interaction networks, and later revitalised by Watts and Strogatz in [WS98], the small world abstraction is essentially that of "small" average path lengths and high clustering coefficients found in most existing networks. More precisely, in a small world network of n nodes the average path length scales as $O(\log n)$ or even slower, also typically with a small constant value hidden by the Landau notation.

Further interesting abstractions have been developed during the past few years. *Scale-free networks* have been discovered as having a ubiquitous role in several areas. In short, scale-free networks have degree distributions with tails following a power law. That is, the probability of randomly selected node having exactly k connections decays as $k^{-\nu}$, in which ν is a positive constant characterising the tail of the distribution further. This is in stark contrast to some earlier used network models, such as random graphs [Bol01], in which the decay is exponential. Several refinements to the original scale-free abstraction have been suggested. For an example arising from Doyle's HOT programme, see [LAWD04]. Additionally, correlations between various graph-theoretic quantities, such as node degrees have given a rise for various abstractions. A recent example is the rich-club phenomenon observed, for example, in Internet topology [ZM04].

Naturally, several *models* have been suggested to construct graphs exhibiting properties characteristic of these abstractions. Most valuable of these models are dynamic in nature, shedding light on the processes that can shape the network to fit to a particular abstraction. The first example is again given by the small-world phenomenon, which can be brought about by *rewiring* some connections in a graph of high clustering coefficient. See figure 1 for an example. Similarly, processes giving rise to other network abstractions have

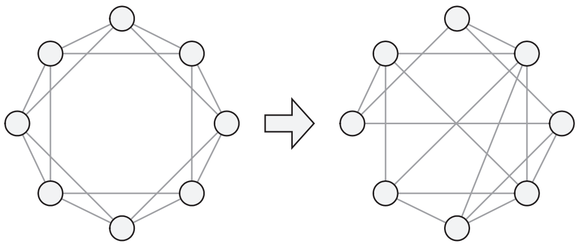

Figure 1: The Watts-Strogatz model of a small-world network based on random rewiring of lattices.

been discovered. Scale-free networks, for example, seem to arise from growth and ageing processes where node lifetimes or attractiveness to new connections is dependent on its degree. These types of processes already give an indication on how abstractions can be exploited in reasoning. Networks can be "coerced" to develop into a particular direction by judicious manipulation of dynamic processes naturally taking place in actively evolving networks. Many of the abstractions encountered have other attractive properties that can likewise be exploited in reasoning processes. For example, scale-free networks tend to be very resilient (in terms of network connectivity) against uniformly random node failures. Thus, from user requirements and desired characteristics, the decision process of selecting the dynamics to follow can be carried out entirely in terms of high-level abstractions, without resorting to the analysis of a particular network model.

We conclude this section by making some comments on wireless networks, and networks with substantial wireless components in them. Most of the work described above is mainly applicable to wired communication networks. Naturally, basic models for wireless networks have been constructed, the most straightforward ones being those based on random geometric graphs [Pen03]. As an example, simple lattice approximation argument shows that wireless networks are certainly *not* small world networks, and are typically expected not to be scale-free, either (this depends on the underlying node location distribution, of course). Basic abstractions in terms of, for example, statistical correlations in node locations, connectivity rules, and, indeed, mobility patterns can be established. Overall, however, understanding which abstractions are relevant for wireless networks, even to the same level of usefulness as the ones developed for the fixed network case, still requires substantial research effort. Especially mobility studies have very much been focussed on developing mobility models for simulation purposes (see [CBD02] for a nice, critical review), as opposed to model-independent characterisation of user mobility.

3 Inference and Reasoning

We shall now briefly discuss the building of the network "self-image", that is, the selection of the abstraction for the present state of the network. If complete topology information is available, the problem is reduced, at least in theory, to simple arithmetic, "goodness-of-fit" testing, and algorithmic decision making. Unfortunately global information is typically *not* available, so local information, possibly coupled with limited probing of the more distant reaches of the network has to be used together with probabilistic inference to yield the abstraction to be used.

For topology probing numerous solutions have been developed. Tools like traceroute, NetInventory [BGJ$^+$04] or Rocketfuel [SMWA04] can be used to discover structures of rather complicated IP-networks, including basic link-layer technologies as well. However, use of such techniques incurs substantial overheads, and requires considerable amount of time. Therefore, methods for inferring network characteristics from a small set of measurements should definitely be studied further. Limited work exists on, for example, fast probabilistic estimation of the average shortest path lengths of network (see [Coh93]), but very little has been done to actually implement these algorithms into concrete protocols. For an interesting approach including probing of path characteristics for QoS maintenance with a testbed implementation, see [GLN04]. For maximum-entropy approach for deriving network models from a limited set of observed data, see [PN04].

Wireless networks pose again further challenges. Since topology of the network is dependent on the underlying geometry, selection of the abstraction should take the geometric relations between the nodes and their surroundings into account. Basic inference problems related to node location distributions have been extensively studied in the context of point processes. For a thorough review on the topic, see [Kar91]. However, the straightforward application of these methods is impossible in most radio networks, as nodes are not aware of their locations. The corresponding localisation problem (without using GPS or similar positioning systems) relying on data fusion techniques has been extensively studied. While techniques such as Bayes filtering seem promising, location information that can be inferred from wireless channels alone is still not very accurate. Thus it might be prudent to strive for more propagation-oriented abstractions for dynamic wireless networks, removing the need for explicit node localization.

4 Conclusions

We have seen that especially for fixed networks a number of powerful network abstractions have been developed, many of which are directly linked to the dynamic processes shaping the networks. These abstractions can be used as a basis for reasoning processes in future networks, to guide on-line topology control, and to optimize the use of network resources. However, several problems have to be solved before protocols of this type become reality. Abstractions for wireless networks should be developed, especially ones capturing the dynamic nature of the users and the wireless environment. Reasoning processes themselves, while intuitive to humans working with these concepts, are not so straightforward to implement on a computer. Before this can take place, we need to develop techniques for representing topology and geometry information in terms of the relevant abstractions. Similarly, we need ways to describe (partial) knowledge on the properties of network dynamics, and their relation to the abstractions used. Solving these problems, and applying the solutions to simulations and real network testbeds for validation should provide for exiting research opportunities far into foreseeable future.

5 Acknowledgments

We are thankful for financial support from DFG (GRK 643: “Software für mobile Kommunikationssysteme”) and RWTH Aachen University. We would also like to thank Konrad Wrona, Marina Petrova and the participants of the Dagstuhl workshop on “Cognitive Networks and Radios” for fruitful discussions.

References

[BGJ$^+$04] Y. Breitbart, M. Garofalakis, B. Jai, C. Martin, R. Rastogi, and A. Silberschatz. Topology discovery in heterogeneous IP networks: the NetInventory system. *IEEE/ACM Transactions on Networking*, 12(3):401–414, June 2004.

[Bol01] B. Bollobas. *Random Graphs*. Cambridge University Press, 2nd edition, 2001.

[CBD02] Tracy Camp, Jeff Boleng, and Vanessa Davies. A Survey of Mobility Models for Ad Hoc Network Research. *Wireless Communication & Mobile Computing*, 2(5):483–502, 2002.

[Coh93] E. Cohen. Fast algorithms for t-spanners and stretch-t paths. In *Proc. 34th IEEE Annual Symposium on Foundations of Computer Science*, pages 648–658, 1993.

[CPRW03] David D. Clark, Craig Partridge, J. Christopher Ramming, and John T. Wroclawski. A knowledge plane for the internet. In *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 3–10, New York, NY, USA, 2003. ACM Press.

[GLN04] E. Gelenbe, R. Lent, and A. Nunez. Self-aware networks and QoS. *Proceedings of the IEEE*, 92(9):1478–1489, September 2004.

[Kar91] A. F. Karr. *Point Processes and their Statistical Inference*. Marcel Dekker, 2nd edition, 1991.

[LAWD04] Lun Li, David Alderson, Walter Willinger, and John Doyle. A first-principles approach to understanding the internet's router-level topology. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 3–14, New York, NY, USA, 2004. ACM Press.

[MRPS04] P. Mähönen, J. Riihijärvi, M. Petrova, and Z. Shelby. Hop-by-hop toward future mobile broadband IP. *IEEE Communications Magazine*, 42(3):138–146, March 2004.

[Pen03] M. Penrose. *Random geometric graphs*. Oxford University Press, 2003.

[PN04] Juyong Park and Mark E. J. Newman. The statistical mechanics of networks. *Phys. Rev. E*, 70, 2004.

[SMWA04] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2–16, February 2004.

[WS98] D. J. Watts and S. S. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.

[ZM04] S. Zhou and R. J. Mondragon. The rich-club phenomenon in the Internet topology. *IEEE Communications letters*, 8(3):180–182, 2004.