Modeling Reality Algorithmically: The Case of Wireless Communication

Magnús M. Halldórsson

ICE-TCS, School of Computer Science, Reykjavik University, 101 Reykjavik, Iceland.
mmh@ru.is

1 Algorithmic Models and Their Properties

Computation is increasingly being viewed as the 21$^{\text{st}}$ century approach to modeling the world. Classical sciences have become increasingly more computer-driven, necessitating a computational perspective. The equation – bastion of 20$^{\text{th}}$ century science – is becoming supplanted by the algorithm. To properly address real-world phenomena, we need models appropriate for algorithmic approaches.

This note contains the author's reflections on the choice and design of models, particularly those capturing aspects of the measurable world. What it is that we look for in models and the essential properties that we seek. We do this in the context of wireless networking, but hope that some of the lessons have wider relevance.

We postulate that algorithmic models must satisfy four properties to be truly useful.

Fidelity A model must be a fair representation of reality. Whereas physics has the advantage that its relatively simple laws hold with extremely high precision, the settings for most fields of study are inherently noisy, making perfect models a pipe dream. Instead, abstractions are an intrinsic part of most models, where the intent is to factor out unimportant ideosyncracies. On the other hand, if essential features are eliminated, the model fails its primary task: to faithfully represent reality.

Simplicity Overly complex models generally result in limited usage. The utility of such models for algorithmic design is necessarily limited, as it complicates all the efforts involved. Occam's razor suggests that models should be as simple as possible, but also no simpler than that, paraphrasing an aphorism attributed to Einstein. Simplicity also has implications for analysis.

Analyzability In order to truly understand real-world phenomena, we need to be able to analyze them and study, both individually and in relation to other phenomena. A model with characteristics that defy analysis may allow for uniformed heuristic use, but will hamper our understanding of the intrinsic properties of the concept at hand.

Generality Finally, we seek explanations of general utility with wide applicability. There is always the danger to introduce context-specific attributes to strengthen the predictive power of the model, but the more we do so, the less useful the model is as a means to explain general properties.

It may be helpful to consider some examples.

Integer Linear Programming ILP is an extremely general tool to represent combinatorial problems, allowing for generic solution methods. Depending on the context, it can be very faithful to the real phenomena. The utilities for algorithm design and analysis depends a lot on the specific domain, and can range from very high to minimal.

Maxwell's equations The equations for electrodynamics that underlie electrical and communication technologies are both very general and highly precise, omitting only the quantum effects that are usually immeasurably small. However, when examined at the scale of wireless networks, the details involved are overwhelming, rendering them unusable for all but exceptional settings of algorithms and analysis.

In general, the utility of a model may depend on the issue/problem under consideration.

2 Selected Wireless Models

A fair number of algorithmic models has been proposed for wireless networks. Let us consider the more prominent models, explore the problems that they address well and examine the issues they raise and their weaknesses. Usually, the distributed setting is assumed, but one can also evaluate them with respect to centralized algorithms. Let n denote the number of wireless transceiver nodes.

Radio model In the earliest and the most basic model, a wireless transmission is successful if exactly one transmitter is transmitting, in which case all the other nodes receive the message.

A core problem addressed in this model, which has been extended to other models, is *leader election*: the nodes should agree on a single node as a leader. With this primitive, many other issues are simplified. One surprising result due to Willard is that this can be achieved in $O(\log \log n)$ steps [15], when the nodes have *collision detection*, i.e., can distinguish silence from the case when two or more nodes are attempting to transmit.

The key limitation of this model is the assumption that all nodes are within communication range. It is also pessimistic in that it does not allow for any spatial reuse of the wireless channel.

General graphs In the (general) graph model, the graph represents which pairs of nodes can communicate (and interfere) with each other. In this sense, the radio model corresponds to the clique graph.

The prototypical problem addressed in this model is the *broadcast* problem: how to transmit a message from a given source to all other nodes in the graph. A celebrated result of Bar-Yehuda, Goldreich and Itai [2] shows that this can be achieved in $O(D \log n)$ time steps with a randomized distributed algorithm, where D is the diameter of the graph. This is essentially optimal for a distributed algorithm, and within logarithmic factor of the best possible by a centralized algorithm.

One downside of the model is that many problems become hard to solve even approximately. For instance, the coloring problem, which captures core questions regarding scheduling wireless communications, has no sublinear approximation algorithm [6].

Disc graphs Communication occurs in the physical world, which is three-dimensional, and distances do matter. A natural approach to limit complexity is therefore to embed the nodes in a Euclidean space and assume that nodes can connect if they are sufficiently close. In the basic setting, nodes can communicate (and interfere) if and only if they are within a fixed distance apart, giving rise to *unit disc graphs* (UDG). UDGs have been the source of a large amount of interesting theory, with the early paper of Clark et al. [3] cited over 1000 times.

Many variations and extensions exist, such as allowing for differing power/radii of the nodes (disc graphs) or different ranges for interference than for communication (protocol model). All disc graphs, however, make strong assumptions: the world is flat (i.e., planar), radio transmission ranges are circular, and reception is symmetric. More generally, all graph-based models assume that reception quality is a binary and that interference is a pairwise relationship. Numerous empirical results (see, e.g., [10]) have shown these to be simplistic.

Physical model The model of choice in engineering circles has been the *physical model*, where the radio signal is assumed to be a decaying function of distance. Here, interference is no longer binary but additive, with successful reception achieved if the total amount of interference is sufficiently small relative to the strength of the intended signal (i.e., high enough signal-to-interference-and-noise-ratio, SINR).

The standard assumption is that signal decays as a polynomial function of distance, known as *geometric pathloss*. Namely, if the signal travels distance d from a sender transmitting with power P, it will be received with strength P/d^α, where α is an absolute constant depending on the setting, understood to be in the range $[2, 6]$.

The physical model was mostly ignored by algorithm theory for a long time, assumed to be too complicated and hard to analyze (failing our Simplicity and Analyzability axioms). Recent years have, however, seen great improvements in our understanding of the model and increase in analytic results.

One problem for which results in the physical model are qualitatively different from those in other models is the *aggregation* problem (and the related connectivity problem): Compute aggregation statistic (say, the minimum) of a set of values, where each wireless node contains only a single value. In any disc or graph-based model, the worst-case round complexity is necessarily linear. Surprisingly, perhaps, Moscibroda and Wattenhofer showed that in the physical model it is only poly-logarithmic [13]. In fact, the worst-case bound is only $O(\log n)$ [7].

Even though the physical model adds several attributes of realism, it still has issues regarding fidelity. Geometric path loss means assuming that antennas are omnidirectional and that signal decays smoothly as a function of distance. Real environments have obstacles and walls that can reflect, scatter and damper signals, and the mere appearance of floors or a ground introduces multi-path effects that are beyond the pure geometric path loss.

3 Future Directions

The preceding models have now all been fairly well studied. Each has aspects that fit certain problems better than others, allowing us to draw distinct lessons. None, however, captures all the important aspects of real environments. Experimental evidence has indeed found that wireless reception is tricky, defying simplistic characterizations [1]. We point out a few additional approaches that have been considered.

If the assumption of geometric path loss is jettisoned from the physical model, we are left with the *abstract SINR* model. This is extremely general, with general graphs being a special case. Thus, scheduling-type problems become highly inapproximable. Still, it may be instructive to consider this general model further, identifying other types of restricted instances or parameterized properties that allow us to recover the Analyzability axiom. The *inductive independence* or *maximum average affectance* property of [9] is one such candidate.

Temporal variability in wireless signal reception has been captured in the recent *dual graph* model [12]. It extends the (general) graphs model by allowing for both reliable and adversarially chosen unreliable links. Whether this exact definition is the best one remains to be seen.

Random artifacts appear to be unavoidable in real networks, at least at a low level. Different versions are known as “fading”, “shadowing”, or “Gaussian” noise. One of the more common ones, Rayleigh fading, has been analyzed in conjunction with the physical model [4] under the assumption of full independence. Correlations and other variations await further study.

Our coverage is by no means exhaustive. Among exciting recent directions are the Abstract MAC Layer [11], multi-channel models (e.g., [5]), jamming resistance [14], and MIMO extensions of the physical model (e.g., [8]).

We believe that the time is ripe for tackling the challenge of faithfully modeling real wireless environments, while obeying the other axioms of simplicity, generality and analyzability. A natural direction would be to meld some of the recent variations with the classical models and assess the resulting model according to these criteria.

Acknowledgements

I thank the attendees of WRAWN (Workshop on Realistic models for Algorithms in Wireless Networks) 2013 for stimulating presentations and discussions that motivated these reflections. Any misrepresentations or omissions are, of course, mine only.

References

1. N. Baccour, A. Koubaa, L. Mottola, M. A. Zuniga, H. Youssef, C. A. Boano, and M. Alves. Radio link quality estimation in wireless sensor networks: a survey. *ACM Transactions on Sensor Networks (TOSN)*, 8(4):34, 2012.
2. R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in radio networks: an exponential gap between determinism and randomization. In *PODC*, pages 98–108, 1987.
3. B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete mathematics*, 86(1):165–177, 1990.
4. J. Dams, M. Hoefer, and T. Kesselheim. Scheduling in wireless networks with Rayleigh-fading interference. In *SPAA*, pages 327–335, 2012.
5. S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Gossiping in a multi-channel radio network. In *DISC*, pages 208–222. Springer, 2007.
6. U. Feige and J. Kilian. Zero knowledge and the chromatic number. In *CCC*, pages 278–287. IEEE, 1996.
7. M. M. Halldórsson and P. Mitra. Wireless Connectivity and Capacity. In *SODA*, 2012.
8. T. Janson and C. Schindelhauer. Broadcasting in logarithmic time for ad hoc network nodes on a line using mimo. In *SPAA*, pages 63–72. ACM, 2013.
9. T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *DISC*, pages 163–178, August 2010.
10. D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *MSWiM*, pages 78–82. ACM, 2004.
11. F. Kuhn, N. Lynch, and C. Newport. The abstract MAC layer. *Distributed Computing*, 24(3-4):187–206, 2011.
12. F. Kuhn, N. Lynch, C. Newport, R. Oshman, and A. Richa. Broadcasting in unreliable radio networks. In *PODC*, pages 336–345. ACM, 2010.
13. T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *INFOCOM*, 2006.
14. A. Ogierman, A. Richa, C. Scheideler, S. Schmid, and J. Zhang. Competitive MAC under Adversarial SINR. *ArXiv e-print 1307.7231*, July 2013.
15. D. E. Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM Journal on Computing*, 15(2):468–477, 1986.