

# Research Statement

Calvin Newport

December 6, 2010

**Research Mission:** To develop distributed algorithms that improve wireless networks.

Distributed algorithms improve systems. We have repeatedly seen these algorithms begin in the theory community before migrating to real systems where they enabled performance breakthroughs. The massive data-centers now driving the Internet, for example, owe much to early theoretical work on consensus. Wireless networking, however, perhaps alone among the major networking technologies, has not yet seen this convergence between distributed algorithms and real systems. Though many theoreticians study distributed algorithms in this setting, their work is largely ignored by practitioners. It is as if the two communities are speaking different languages. This is a problem: by keeping these two concerns separate, we are potentially missing out on novel wireless system designs that would enable more functionality and better performance. I want to solve this problem by tackling hard theory problems that help real networks perform better.

Here are two reasons to believe I can fulfill my mission: First, my training. I received my PhD in a theory group and am currently a postdoctoral associate in a systems group. Put another way, I am a theoretician, but I also speak the language of the networking community and actively seek collaboration. The second reason is my research record. A major goal of my work is to solve hard theory problems that are relevant to real networks. As summarized in Section 1, my efforts have resulted in a large number of publications in top conferences and journals [6, 9, 8, 39, 7, 16, 17, 20, 15, 33, 34, 40, 41, 2], and have spawned follow up work from both the theory (e.g., [38, 28, 10, 13, 11, 29]) and systems (e.g., [51, 50, 42, 43]) communities. On becoming a postdoc I adopted a three-part research strategy: (1) collaborate with wireless network researchers to identify real problems that might be solved by smart distributed algorithms; (2) solve the identified theory problems; (3) when possible, experimentally evaluate the algorithms' benefits. In Section 2, I describe five on-going projects generated by this strategy.

## 1 Previous Work

As a graduate student, I studied distributed algorithm problems that were both *interesting* from a theory perspective and *practical* from a systems perspective. In this section, I summarize the results of these efforts.

A major topic investigated by my research is realistic wireless network models. My collaborators and I observed that most theoretical work on algorithms in this setting simplify communication into a series of deterministic rules. The correctness of algorithms proved in these models often depends on the inviolable nature of these rules.<sup>1</sup> In practice, of course, wireless communication is anything but predictable. *This*

---

<sup>1</sup>A classic example is the problem of broadcasting a message through a constant-diameter network. If you assume it is possible for a process to *sometimes* receive one message from among several during a collision,  $\Omega(n)$  steps are required to solve the problem [3, 4]. If you assume that *all* colliding messages are *always* lost, however, you can solve the problem in  $\mathcal{O}(\log n)$  steps [30].

**presents a problem:** Protocols proved correct in the classic models with deterministic receive rules might fail when deployed in the chaotic world of real wireless networks. At the same time, however, if we modify models to allow for unpredictable message loss behavior, it might be difficult to prove *any* useful results. In a series of papers I co-authored while a graduate student [6, 9, 8, 9, 39, 21, 7, 16, 17, 22, 20, 15], we addressed this problem by demonstrating—perhaps surprisingly—that that even in difficult settings with unpredictable, correlated, and even adversarial message disruption, it is *still* possible to design efficient solutions to fundamental coordination problems. We argued that these pessimistic models should be used to develop critical *control protocols*, which are run infrequently, but whose correctness is crucial for good system performance. Below, I briefly summarize the two models we studied and the results we proved.<sup>2</sup>

**The Disrupted Radio Network Model.** In [16], we introduced the *disrupted radio network* (DRN) model, which includes the *interference adversary* formalism. During each time step, this adversary can disrupt a bounded subset of the available communication channels. The interference adversary does not represent a literal adversarial device, but instead incarnates the diversity of unpredictable disruption behaviors that occur in real networks. In a series of papers published at PODC [17, 15], DISC [16], and INFOCOM [20], we present a collection of techniques for solving the fundamental problem of information exchange in this adversarial model. In early work [16], drawing on an unexpected connection to Turán’s Theorem [52] (a seminal result from extremal graph theory), we proved that *oblivious* deterministic solutions to this problem are inherently slow. We then proved the surprising result that *adaptive* deterministic solutions to information exchange, by contrast, can perform within a constant factor of optimal, assuming that the number of disrupted channels at every time step is bounded by the root of the total number of channels [20]. In [17], we produced efficient randomized solutions to this problem, and then used them to establish a group key agreement protocol. And in [15], we explored the related problem of agreeing on a value—e.g., a leader id, round labeling, or frequency hopping pattern seed—among an unknown number of devices activated during different rounds. In my dissertation [40], among other results, I studied additional problems in this model, including reliable broadcast. Emphasizing the practical nature of this work is the fact that the model has been adopted by a growing number of systems researchers, at venues such as MobiHoc [50, 42], the IEEE Symposium on Security and Privacy [51], and the Journal on Selected Areas in Communications [43].

**The Unreliable Collision Detection Model.** In a series of papers published at PODC [9], Allerton [8], ASDN [6], and the journal *Distributed Computing* [7], we studied the fundamental problem of consensus in a single-channel variant of the DRN model. In this model, we assume that each receiver can lose an arbitrary subset of transmitted messages in each round. To circumvent obvious impossibilities, we also introduce the *unreliable collision detector* formalism, which attempts to provide each receiver with a binary indication of whether or not *some* transmitted messages were lost. Among other results, we provide a constant-time consensus solution for a collision detector that notices when half or more of the transmitted messages in a round are lost, and show a tight logarithmic bound for weaker detectors [9, 6, 7]. (These results are also expanded and detailed in my masters thesis [39].) We then explored the use of unreliable detectors to build a general middleware layer to facilitate reliable coordination in wireless networks [8]. Subsequent follow-up work by systems researchers validated that the most basic detectors studied in our work are easily implemented using the on-chip *clear channel assessment* algorithm integrated with the popular CC2420 radio [14].

Another major focus of my research is the study of *abstraction layers* that split the task of designing

---

<sup>2</sup>These models describe the network from the perspective of the *link layer*. Instead of attempting to describe the complicated physical layer behaviors causing packets to be unexpectedly lost (as is done, for example, with the SINR model), we describe only the effect of these behaviors on the packets themselves. We also omit details such as bit rate, as the style of control protocol motivating this work focus on correctness, not throughput.

wireless algorithms into two sub-tasks: (1) proving an algorithm correct when combined with the (presumably, easy to use) abstraction; and (2) implementing the abstraction in realistic models. Most theoreticians can concentrate on producing distributed algorithms that use the abstractions, while only a masochistic few need tackle the messy task of implementing these abstractions. Below, I summarize my two projects concerned with these goals.

**The Abstract MAC Layer.** In a paper presented at DISC [33] and subsequently in the journal *Distributed Computing* [34], Fabian Kuhn, Nancy Lynch, and I, introduced the *abstract MAC layer* formalism. This abstraction provides a powerful reliable local broadcast service, but can also be implemented in a variety of realistic models. We validate the abstraction’s usefulness by using it to produce new results for the problem of *multi-message broadcast*. By allowing us to separate low-level concerns regarding contention from high-level message pipelining issues, the abstraction helped us obtain strong results for a problem that had been previously neglected due to its complexity. Several other research teams have since embraced this formalism to achieve similar breakthroughs [28, 10, 13, 11, 29].

**A General Wireless Network Composition Framework.** In my dissertation [40], a CONCUR paper [41], and an invited journal paper currently under submission to *Distributed Computing*, Nancy Lynch and I explored a more general abstraction strategy for wireless networks. We defined a formal model for probabilistic algorithms communicating on an arbitrary radio network (each component is modeled as a probabilistic automaton). We then defined what it means to *implement* one network type with another, and prove *composition* results that show that an algorithm that solves a problem in one type of network will still solve the problem when combined with an implementation of that network. This work generalizes the abstract MAC layer—allowing algorithm designers to first prove their algorithms correct in a simple and easy to use network model, and then later implement the easy model in a more realistic setting.

My final major wireless networking project concerned fundamental limits on a pervasive style of attack.

**Fundamental Limits for Denial of Service Attacks.** Because radio waves are an open medium, wireless networks are vulnerable to jamming-style denial of service attacks. Presumably, an adversary wants to avoid continuous jamming, as this attack is easy to detect (and subsequently eliminate). A better strategy is to jam *just enough* key messages to prevent a protocol from terminating. Previous work studied the achievable efficiency of these attacks for specific protocols, such as the 802.11 DCF [5]. By contrast, in a conference paper at OPODIS [21], and a journal paper in *Theoretical Computer Science* [22], Seth Gilbert, Rachid Guerraoui and I, studied a more fundamental question: what is the best possible jamming efficiency that can be achieved against *all* protocols. In more detail, we proved that for a wide collection of problems, including single message broadcast, consensus, and leader election, for every protocol that solves one of these problems, there is a denial of service strategy that requires the adversary to jam no more than half of the messages. We then provide matching upper bounds. In recent work published at SPAA [2], working now also with Dan Alistarh and Zarko Milosevic, we expand some of these results to a multihop setting.

## 2 New Work

In the fall of 2009, I started as a postdoctoral associate working with Hari Balakrishnan in MIT’s Networks and Mobile Systems Group. I treated the first two semesters of my postdoc as a crash course in systems research. During this time, I conducted a series of vehicular network measurement studies, the results of which I presented in a technical talk at Ford Motor Company’s Dearborn headquarters. I also collaborated

on a project to use mobile device sensors to improve the performance of low-level wireless protocols, the results of which are presented in a HotNets [44] and NSDI [45] paper.

My apprenticeship complete, I turned my attention to my **three-part research strategy**: (1) collaborate with wireless network researchers to identify real problems that might be solved by distributed algorithms; (2) work through the necessary theory for the algorithms; (3) when possible, experimentally evaluate the algorithms' benefits. Using this strategy, I have launched five projects, which I summarize below. Before continuing with the summaries, however, I want to note that the first four projects all concern a new research direction that I call **subversive networking**. This direction explores the use of distributed coordination among wireless devices to improve the performance of infrastructure networks (e.g., WiFi or 3G).<sup>3</sup> This improvement can be in terms of throughput to the infrastructure, but can also be in terms of minimizing battery or cellular bandwidth usage.

**Distributed Intelligence in Vehicular Networks.** Vehicles in an urban setting can make local observations such as traffic conditions and accident alerts, that when disseminated and combined can help users make intelligent use of the city's transportation systems. (See, for example, MIT's CarTel Project [23], which motivated this work.) In a recent paper published at POMC [12], Alejandro Cornejo and I present a distributed algorithm for maintaining such a system using only local communication. True to the subversive networking philosophy, this approach saves the expense (in terms of cellular bandwidth costs and battery power) of requiring every user to continuously send observations over cellular links. This local approach also scales better than centralized solutions. Here are three interesting notes regarding this work: (1) unlike most theoretical studies of mobile networks, we used extensive position traces from taxicabs in Boston (obtained from CarTel) to construct and validate our model; (2) our algorithm's properties depend in interesting ways on optimal solutions to the classic *Towers of Hanoi* problem [47, 31]; and (3) researchers from Ford Motor Company's Infotronics Research and Advanced Engineering team are implementing our algorithm in simulation, with an eye towards potential test bed deployment. (See this article [37] from *MIT News* for more about this project and Ford's interest.)

**Reducing Cellular Bandwidth Usage With Distributed Aggregation.** The Comfort Cab Company in Singapore maintains a fleet of 15000 taxi cabs. The company wants to monitor the occupancy counts in different regions of the city. Assuming that each cab reports once every 5 minutes, and each interaction with the server requires 256 bytes, the fleet will consume over 550 GB of bandwidth in a single 12 hour period. The expense of bandwidth (there are no unlimited 3G plans available), and the infrastructure required to handle this volume of connections, is non-trivial. Motivated by examples of this type, I am studying, working with collaborators at MIT and the National University of Singapore, algorithms that minimize the cellular connections required to aggregate data in mobile networks. Our algorithms use local dissemination to prevent every user from needing to report to the central server. The algorithms are provably *lossless*: the server learns the same information as if every user reported individually. In addition, using techniques from *dynamic graph theory* [27, 26, 46, 24, 36], we can bound the number of cellular connections by the number of "sufficiently dense" regions in the changing communication graph. At a high-level: the algorithm takes advantage of every possible opportunity to safely aggregate data. We are currently preparing a theory paper that describes the techniques we developed for designing this style of algorithm. Furthermore, Seth Gilbert, a professor at the National University of Singapore, and I, recently submitted an invited proposal to the University's Future of Urban Mobility project, to evaluate the benefits of our algorithms using the massive position trace database they obtained from the Comfort Cab Company. If the evaluation goes well, there is the potential to deploy these algorithms in their cabs.

---

<sup>3</sup>The name refers to the fact that devices *subvert* the authoritarian rule of the infrastructure by performing some collusion among themselves.

**Improving Cellular Performance Using WiFi.** WiFi connections offer two advantages over cellular connections: they require less energy and are typically higher bandwidth. It makes sense, therefore, for mobile devices to augment (or replace) their 3G connections with WiFi connections, when available. Working with other members of the Networks and Mobile Systems Group at MIT, I am currently studying this problem for vehicular networks. Whereas previous work focused on maximizing the duration of 1-hop links between vehicles and nearby WiFi access points [19], we are studying how cars can cooperate to establish *short multihop routes* to WiFi infrastructure—increasing the availability of this access. In work published at HotNets [44] and NSDI [45], we presented mesh route metrics that make use of GPS compass hints to help select for long-lived links in a vehicular context. We have shown in simulation, using real vehicle position traces, that these metrics can increase the stability of routes by a factor of 3 to 4. I am currently co-supervising a masters student who is implementing and evaluating these strategies in simulation.

**The Power of Signaling in Wireless Networks.** Over the past two years, several researchers have stumbled onto the same idea: in many wireless network settings, coordinating clients can improve performance (for example, making smarter decisions about which client should go next in uploading packets to an access point) [53, 48, 18]. Such coordination, however, is usually impractical: the overhead of the coordination packets swamp the improvement generated from better coordination. These researchers continue by noting that it is possible to modify existing radios to allow nearby devices to *signal* each other without disrupting the main data traffic. These signals can be used to coordinate devices in a manner that increases overall system performance. Different strategies are proposed to implement these signals, from adding detectable interference patterns to DSSS modulated 802.15.4 data packets [53], to broadcasting detectable known sequences at low SNR [18]. The coordination protocols proposed in [53, 48, 18], however, are basic, and this leaves open an intriguing cross-discipline question: *What is possible in this signaling context with more powerful distributed algorithms?*

Motivated by this question, I developed two signaling-based distributed algorithms. The first is an efficient weighted leader election scheme that selects for the client with the largest packet queue. This algorithm can be run concurrently with the data traffic at an access point to select the next uplink client. The second algorithm provides rich acknowledgments for a sender broadcasting streaming media to an unknown and changing set of receivers. It allows the sender to calculate, concurrently with its transmissions, a distribution of the loss rates experienced by the receivers—enabling nuanced adjustments to its coding and bit rate. I am currently working with collaborators at MIT to experimentally evaluate the benefits of such algorithms.

**Optimal Algorithms for Mesh Networks.** The spread of mesh networks has lagged behind the optimistic predictions that accompanied the rise of 802.11. There are many reasons for this lag, most of which center on the unexpected difficulties of deploying these networks. Of these difficulties, perhaps the most prominent is varying link quality [1]. Not only do wireless links have different qualities, but these qualities change over time (sometimes quite rapidly [49]). This presents a problem because it is not obvious how to best construct and repair structures in mesh networks with unpredictable link behavior. In work published at the last two PODC conferences [32, 35], working in collaboration with researchers from MIT, Arizona State University, and the University of Lugano, we introduced a theoretical network model that captures this issue of unpredictable link behavior, and studied algorithmic strategies for compensating for these difficulties. In new work currently being prepared for submission, we use this model to study the specific problem of building and repairing useful structures in mesh networks. Among other results, we show that the achievable efficiency of these algorithms depends on the accuracy of link quality estimations. We also provide tight upper bounds. The next step for this work is to seek collaboration with a mesh network researcher to evaluate whether our theoretically-derived strategies offer benefits in practice.

### 3 Future Work

The previous section describes a series of on-going projects that will occupy much of my research attention for the next year. Prognosticating beyond this point is difficult as my research strategy leverages collaboration with systems researchers to identify interesting problems. *What* I am working on two to three years from now will depend heavily on *who* I am working with. Therefore, instead of attempting to lay out a rigid vision of where my work is headed, I conclude this research statement with a rapid-fire list of the ideas currently in the *on deck* circle of my research queue—providing a sense of the type of problems that interest me. I fully expect this list to mutate and grow.

- **An Efficient Mobile Sensor Network Platform.** A hot research topic is the use of smart phones to form mobile sensor networks. A major concern with such networks is reducing battery and cellular bandwidth usage. I want to design a mobile sensor network platform that uses distributed coordination and aggregation among nearby devices to reduce redundant sensing and minimize cellular reports. In my vision, users could set an energy and bandwidth budget, and the platform would maximize the useful data returned given these constraints.
- **Predicting the Future in Mobile Networks.** There exists a fascinating series of unexpected mathematics results on the topic of *predicting the future*; c.f., [25]. (In slightly more detail, given an adversary that controls a set of variables that each evaluate to either *good* or *bad* in each round, they show how to use past observations to choose good variables at roughly the same rate as if you know the outcomes in advance.) Working with a researcher at EPFL, I have been investigating the application of these results to predicting long-lived pairwise links in mobile networks—with the goal of maximizing connectivity.
- **Ultra-Low Power Configuration Using Signaling.** In the previous section, I discussed on-going work to study distributed algorithms that leverage *signaling* to coordinate devices. I am interested in exploring ultra-low power implementations of this signaling (based, for example, on RFID resonator technology), and distributed algorithms that use these schemes to coordinate energy-constrained networks; e.g., allowing a device to keep its 802.11 radio circuits powered down until its next turn to send or receive data.
- **Toward a Wireless Network Model Hierarchy.** The theoretical literature on wireless networks abounds with a diversity of network models. Results in one model are typically incompatible with results in another. It seems, however, as if most models can be described by two factors: how much information devices have about the network topology, and at what rate can they learn new information about the topology. I want to make this observation explicit: unifying these existing models in a canonical *topology oracle* model, and therefore allowing formal comparisons of these models' power.

### References

- [1] D. Aguayo, J. Bicket, S. Biswas, R. Morris, B. Chambers, and D. De Couto. MIT Roofnet. In *Proceedings of the International Conference on Mobile Computing and Networking*, 2003.
- [2] D. Alistarh, S. Gilbert, R. Guerraoui, Z. Milosevic, and C. Newport. Securing Every Bit: Authenticated Broadcast in Radio Networks. In *Proceedings of the Symposium on Parallel Algorithms and Architectures*, 2010.
- [3] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45(1):104–126, 1992.
- [4] R. Bar-Yehuda, O. Goldreich, and A. Itai. Errata regarding "on the time complexity of broadcast in radio networks: An exponential gap between determinism and randomization". [http://www.wisdom.weizmann.ac.il/mathusersodedp\\_bgi.html](http://www.wisdom.weizmann.ac.il/mathusersodedp_bgi.html), 2002.

- [5] T. X. Brown, J. E. James, and A. Sethi. Jamming and sensing of encrypted wireless ad hoc networks. Technical Report CU-CS-1005-06, UC Boulder, 2006.
- [6] G. Chockler, M. Demirbas, S. Gilbert, N. Lynch, C. Newport, and T. Nolte. Reconciling the theory and practice of unreliable wireless broadcast. In *International Workshop on Assurance in Distributed Systems and Networks (ADSN)*, June 2005.
- [7] G. Chockler, M. Demirbas, S. Gilbert, N. Lynch, C. Newport, and T. Nolte. Consensus and collision detectors in radio networks. *Distributed Computing*, 21:55–84, 2008.
- [8] G. Chockler, M. Demirbas, S. Gilbert, and C. Newport. A middleware framework for robust applications in wireless ad hoc networks. In *Proceedings of the 43rd Allerton Conference on Communication, Control, and Computing*, 2005.
- [9] G. Chockler, M. Demirbas, S. Gilbert, C. Newport, and T. Nolte. Consensus and collision detectors in wireless ad hoc networks. In *Proceedings of the International Symposium on Principles of Distributed Computing*, pages 197–206, New York, NY, USA, 2005. ACM Press.
- [10] H. Chung, P. Robinson, and J. Welch. Regional Consecutive Leader Election in Mobile Ad-Hoc Networks. In *Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, 2010.
- [11] A. Cornejo, N. Lynch, S. Vigar, and J. Welch. Neighbor Discovery in Mobile Ad Hoc Networks Using an Abstract MAC Layer. In *Proceedings of the Annual Allerton Conference on Communication, Control*, 2010.
- [12] A. Cornejo and C. Newport. Prioritized Gossip in Vehicular Networks. In *Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, 2010.
- [13] A. Cornejo, S. Vigar, and J. Welch. Reliable Neighbor Discovery for Mobile Ad Hoc Networks. In *Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, 2010.
- [14] M. Demirbas, O. Soysal, and M. Hussain. A singlehop collaborative feedback primitive for wireless sensor networks. *Proceedings of the Conference on Computer Communication*, 2008.
- [15] S. Dolev, S. Gilbert, R. Guerraoui, F. Kuhn, and C. Newport. The wireless synchronization problem. In *Proceedings of the International Symposium on Principles of Distributed Computing*, 2009.
- [16] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Gossiping in a multi-channel radio network: An oblivious approach to coping with malicious interference. In *Proceedings of the International Symposium on Distributed Computing*, 2007.
- [17] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Secure Communication Over Radio Channels. In *Proceedings of the International Symposium on Principles of Distributed Computing*, 2008.
- [18] A. Dutta, D. Saha, D. Grunwald, and D. Sicker. SMACK: a SMART ACKnowledgment scheme for broadcast messages in wireless networks. *ACM SIGCOMM Computer Communication Review*, 39(4):15–26, 2009.
- [19] J. Eriksson, H. Balakrishnan, and S. Madden. Cabernet: Vehicular Content Delivery using WiFi. In *Proceedings of the International Conference on Mobile Computing and Networking*, 2008.
- [20] S. Gilbert, R. Guerraoui, D. Kowalski, and C. Newport. Interference-Resilient Information Exchange. In *Proceedings of the Conference on Computer Communication*, 2009.
- [21] S. Gilbert, R. Guerraoui, and C. Newport. Of malicious motes and suspicious sensors: On the efficiency of malicious interference in wireless networks. In *Proceedings of the International Conference on Principles of Distributed Systems*, December 2006.
- [22] S. Gilbert, R. Guerraoui, and C. Newport. Of malicious motes and suspicious sensors: On the efficiency of malicious interference in wireless networks. *Theoretical Computer Science*, 410(6-7):546–569, 2009.
- [23] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: a Distributed Mobile Sensor Computing System. In *Proceedings of the Conference on Embedded Networked Sensor System*, 2006.
- [24] M. Jelasity, A. Montessor, and O. Babaoglu. Gossip-Based Aggregation in Large Dynamic Networks. *ACM Transactions on Computer Systems*, 23(3):219–252, 2005.
- [25] A. Kalai and S. Vempala. Efficient Algorithms for Online Decision Problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [26] D. Kempe and J. Kleinberg. Protocols and Impossibility Results for Gossip-Based Communication Mechanisms. In *Proceedings of the Symposium on Foundations of Computer Science*, 2002.
- [27] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and Inference Problems for Temporal Networks. In *Proceedings of the Symposium on Theory of Computing*, 2000.
- [28] M. Khabbazian, F. Kuhn, D. Kowalski, and N. Lynch. Decomposing Broadcast Algorithms using Abstract MAC Layers. In *Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, 2010.

- [29] M. Khabbazi, F. Kuhn, N. Lynch, M. Medard, and A. ParandehGheibi. MAC Design for Analog Network Coding. *MIT CSAIL Technical Report, MIT-CSAIL-TR-2010-036*, 2010.
- [30] D. R. Kowalski and A. Pelc. Deterministic broadcasting time in radio networks of unknown topology. In *Proceedings of the Symposium on Foundations of Computer Science*, pages 63–72, 2002.
- [31] M. Kraitchik. *Mathematical Recreations*, chapter The Tower of Hanoi, pages 91–93. W. W. Norton, 1942.
- [32] F. Kuhn, N. Lynch, and C. Newport. Brief announcement: Hardness of broadcasting in wireless networks with unreliable communication. In *Proceedings of the International Symposium on Principles of Distributed Computing*, 2009.
- [33] F. Kuhn, N. Lynch, and C. Newport. The Abstract MAC Layer. In *Proceedings of the International Symposium on Distributed Computing*, 2009.
- [34] F. Kuhn, N. Lynch, and C. Newport. The Abstract MAC Layer. *Distributed Computing*, To Appear.
- [35] F. Kuhn, N. Lynch, C. Newport, R. Oshman, and A. Richa. Broadcasting in Unreliable Radio Networks. In *Proceedings of the International Symposium on Principles of Distributed Computing*, 2010.
- [36] F. Kuhn, N. Lynch, and R. Oshman. Distributed Computation in Dynamic Networks. In *Proceedings of the Symposium on Theory of Computing*, 2010.
- [37] M. N. O. Larry Hardesty. Cars as Traffic Sensors, September 2010.
- [38] D. Meier, Y. A. Pignolet, S. Schmid, and R. Wattenhofer. Speed dating despite jammers. In *The Proceedings of the International Conference on Distributed Computing in Sensor Systems*, 2009.
- [39] C. Newport. Consensus and collision detectors in wireless ad hoc networks. Master’s thesis, MIT, 2006.
- [40] C. Newport. *Distributed Computation on Unreliable Radio Channels*. PhD thesis, MIT, 2009.
- [41] C. Newport and N. Lynch. Modeling radio networks. In *Proceedings of the International Conference on Concurrency Theory*, 2009.
- [42] G. Noubir, T. Jin, and B. Thapa. Zero pre-shared secret key establishment in the presence of jammers. In *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing*, 2009.
- [43] C. Popper, M. Strasser, and S. Capkun. Anti-Jamming Broadcast Communication Using Uncoordinated Spread Spectrum Techniques. *The IEEE Journal on Selected Areas in Communications*, 28(5):703–715, 2010.
- [44] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden. ‘Extra-Sensory Perception’ for Wireless Networks. In *Proceedings of the ACM Workshop on Hot Topics in Networks*, 2010.
- [45] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden. Improving Wireless Network Performance Using Sensor Hints. In *Proceedings of the Symposium on Networked Systems Design and Implementation*, 2011.
- [46] C. Scheideler. Models and Techniques for Communication in Dynamic Networks. In *Proceedings of the International Symposium on Theoretical Aspects of Computer Science*, 2002.
- [47] P. H. Schoute. De Ringen van Brahma. *Eigen Haard*, 22:274–276, 1884.
- [48] S. Sen, N. Santhapuri, R. Choudhury, and S. Nelakuditi. Moving Away from Collision Avoidance: Towards Collision Detection in Wireless Networks. 2009.
- [49] K. Srinivasan, M. Kazandjieva, S. Agarwal, and P. Levis. The  $\beta$ -Factor: Measuring Wireless Link Burstiness. In *Proceedings of the Conference on Embedded Networked Sensor System*, 2008.
- [50] M. Strasser, C. Pöpper, and S. Capkun. Efficient uncoordinated fhss anti-jamming communication. In *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing*, 2009.
- [51] M. Strasser, C. Pöpper, S. Capkun, and M. Cagalj. Jamming-resistant key establishment using uncoordinated frequency hopping. In *The IEEE Symposium on Security and Privacy*, 2008.
- [52] P. Turán. On an extremal problem in graph theory. *Matematicko Fizicki Lapok*, 48, 1941.
- [53] K. Wu, H. Tan, Y. Liu, J. Zhang, Q. Zhang, and L. M.Ni. Free Side Channel: Bits Over Interference. In *Proceedings of the International Conference on Mobile Computing and Networking*, 2010.