

Protocol Design Beyond Graph-Based Models

Thomas Moscibroda
Microsoft Research
moscitho@microsoft.com

Roger Wattenhofer
ETH Zurich
wattenhofer@tik.ee.ethz.ch

Yves Weber
ETH Zurich
webery@tik.ee.ethz.ch

ABSTRACT

In this paper we shed new light on the fundamental gap between graph-based models used by protocol designers and fading channel models used by communication theorists in wireless networks. We experimentally demonstrate that graph-based models capture real-world phenomena inadequately. Consequentially, we advocate studying models beyond graphs even for protocol-design. In the main part of the paper we present an archetypal multi-hop situation. We show that the theoretical limits of any protocol which obeys the laws of graph-based models can be broken by a protocol explicitly defined for the physical model. Finally, we discuss possible applications, from data gathering to media access control.

1 INTRODUCTION

Wireless multi-hop networks such as sensor, ad hoc, or mesh networks are often modeled by means of graphs. In the most general model, two graphs are given: A connectivity graph $G_c = (V, E_c)$ and an interference graph $G_i = (V, E_i)$. Both graphs are based on the set of devices V . A receiver v successfully decodes a message from a sender u , if and only if u and v are neighbors in the connectivity graph, $(u, v) \in E_c$, and v does not have a concurrently transmitting neighbor node in the interference graph G_i . Protocol designers often consider special cases of this general model. For example, it is sometimes assumed that $G_i = G_c$, or that G_i is G_c augmented with all edges between 2-hop neighbors in G_c .¹ In graph-based models, a protocol designer has to take care that no neighbor of v in G_i is transmitting simultaneously to a neighbor in G_c , or at least, that this happens rarely.

Graph-based models have been particularly popular with higher-layer protocol designers, as they abstract away real-world complications. On the other hand, the concept of an *edge* is oversimplifying starkly, as it is a binary representation for continuous (non-binary!) physical laws. In fact, nodes barely outside the interference range of a receiver v (that is, a node not connected by an

edge with v in G_i) might still cause enough cumulated interference such that receiver v is not able to decode a message from a legitimate neighboring sender in G_c .

Communication theorists on the other hand often do not employ graph-based models. Instead they are studying an arsenal of fading channel models, the simplest being the signal-to-interference-plus-noise ratio (SINR) model. In the SINR model, the energy of a signal fades with the distance to the power of the path-loss parameter α . If the signal strength received by a device divided by the interfering strength of competitor transmitters (plus the noise) is above some threshold β , the receiver can decode the message, otherwise it cannot. This simple SINR model is unrealistic as well, mostly because it is inherently geometric: In reality antennas are not perfectly isotropic, and even more importantly the environment is obstructed by walls or plants. Although these issues can be integrated into the basic SINR model, these “SINR+” models are predisposed to get complicated – essentially intractable from the point of view of a protocol designer. Graph-based models on the other hand automatically incorporate both imperfect (or even directional) antennas and terrains with obstructions. It seems that a majority of classes, books, or tutorials therefore prefers to teach higher-layer concepts in wireless multi-hop networking in terms of graphs, not in terms of SINR.

Even though SINR models allow for exciting scaling law studies (e.g. the theoretical capacity of wireless multi-hop networks), they are often too complicated to comprehend a protocol, let alone analytically prove correctness and/or efficiency of a protocol.

We believe that bridging the gap between protocol designers and communication theorists is a fundamental challenge of the coming years, a hot topic for the wireless multi-hop community with implications for both theory and practice. In particular, in this paper, we advocate studying models beyond graphs, especially for *protocol-design*. After some introductory back-of-the-envelope calculations in Section 2, Section 3 presents experimental results that show that even vanilla sensor radios with restricted hardware can achieve communication patterns which are impossible in graph-based models. In Section 4 we head beyond these straight-forward examples and fantasize about the applications of the experimental findings; in particular we present an archetypal multi-hop

¹Alternatively, it is sometimes assumed that these graphs are the result of a geometric setting. In particular the nodes are points in the Euclidean plane. Two nodes are neighbors in G_c if their Euclidean distance is at most 1. Two nodes are neighbors in G_i if their Euclidean distance is at most r , for some parameter $r \geq 1$. This model is widely known as the *unit disk graph* model with interference.

situation where we propose routing/transport schemes which may break the theoretical throughput limits of any protocol which obeys the laws of a graph-based model. Sections 5 and 6 then discuss related work and future directions, respectively. Note that our examples are preliminary in the sense that they are geared towards illustrating basic concepts and highlighting the fundamental problems of graph-based modeling, rather than towards maximizing the achievable throughput.

2 MOTIVATING EXAMPLES

Consider a network of n devices x_1, x_2, \dots, x_n . A message from a transmitter x_s can be correctly decoded by a receiver x_r if and only if $\frac{P_r}{I_r+N} \geq \beta$ for a hardware-dependant ratio β . In this equation, P_r is the signal strength of the message at the receiver, I_r is the sum of all interferences at x_r , and N is the ambient noise.

In the *physical model* of signal propagation [8] the signal strength P_r is modeled as a polynomially decreasing function depending of the distance $d(x_s, x_r)$ between the sender and the receiver. More precisely, it is assumed that $P_r = \frac{1}{d(x_s, x_r)^\alpha}$ where α is called the *path-loss exponent*, a constant dependent on the medium, typically between 2 and 6.

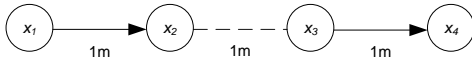


Figure 1: Four nodes placed equidistantly in a line.

Consider the simple “hidden-terminal” network with 4 nodes illustrated in Figure 1.² Nodes x_1 and x_3 want to send a message to the corresponding receivers x_2 and x_4 , respectively. A graph-based communication model implies that at least *two* time slots are required. Otherwise, the two messages would collide at x_2 .

In the physical SINR model, however, the two messages can be easily sent in parallel. For a simple calculation, assume $\alpha = 3$, $\beta = 3$, and background noise $N = 10$ nW. Those values are realistic, even pessimistic, in sensor networks [12] as well as other forms of wireless networks. Let $\beta_{x_i}(x_j)$ be the SINR ratio at a node x_i in which the signal power from node x_j is considered “signal” and the signal power of all other simultaneously transmitting nodes is considered interference. That is, a node x_i successfully receives a message from x_j if and only if $\beta_{x_i}(x_j) \geq \beta$.

If x_1 and x_3 send with power $P_{x_1} = 0$ dBm and $P_{x_3} = -7$ dBm, respectively, we get the following SINR values: $\beta_{x_2}(x_1) = \frac{1000 \mu\text{W}/(1 \text{ m})^3}{0.01 \mu\text{W} + (200 \mu\text{W}/(1 \text{ m})^3)} \approx 5.00$ and $\beta_{x_4}(x_3) =$

²Depending on the specific application scenario, all four nodes may be sensors in a wireless sensor network or stations in a wireless mesh network. Alternatively, nodes x_2 and x_4 may be base stations and x_1 and x_3 may be clients in a Wireless LAN.

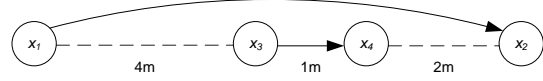


Figure 2: A more elaborate example with four nodes.

$\frac{200 \mu\text{W}/(1 \text{ m})^3}{0.01 \mu\text{W} + (1000 \mu\text{W}/(3 \text{ m})^3)} \approx 5.40$. Consequently, both receivers can correctly decode their corresponding message without any problems.

Now we consider a more elaborate example by rearranging the two sender-receiver pairs (x_1, x_2) and (x_3, x_4) in a way that one pair is placed in the transmission line of the other. This setup is shown in Figure 2. As before, the question is whether it is really necessary to schedule the two messages in succession or if they can be sent in the same time slot without colliding at any of the two receivers. Clearly, any graph-based approach trying to send the two messages in parallel will fail because, intuitively, the medium between x_3 and x_4 can only be used once per time slot.

In the SINR model, however, both messages can easily be transmitted simultaneously, thereby doubling the achieved throughput. When x_1 sends with $P_{x_1} = 1$ dBm and x_3 with $P_{x_3} = -15$ dBm, we get the signal-to-noise and interference ratios of $\beta_{x_2}(x_1) = \frac{1.26 \text{ mW}/(7 \text{ m})^3}{0.01 \mu\text{W} + (31.6 \mu\text{W}/(3 \text{ m})^3)} \approx 3.11$ and $\beta_{x_4}(x_3) = \frac{31.6 \mu\text{W}/(1 \text{ m})^3}{0.01 \mu\text{W} + (1.26 \text{ mW}/(5 \text{ m})^3)} \approx 3.13$. That is, the SINR ratios are such that node x_4 can perfectly decode x_3 ’s message, and at the same time, x_2 successfully receives x_1 ’s message. There is no collision.

3 PROMISING EXPERIMENTS

After these theoretical considerations, this section shows that the effects described above are not only theoretical shenanigan, but can be verified with widely used standard sensor nodes. We decided to employ the mica2 sensor nodes running with TinyOS. They are equipped with a ChipCon CC1000 radio transceiver configured to send at a frequency of 868 MHz.

3.1 Two Pairs of Nodes

We created a testbed with two senders x_1 and x_3 and two corresponding receivers x_2 and x_4 positioned on a line similar to the setup shown in Figure 2. The distances between the nodes were scaled down to 100 cm, 30 cm, and 60 cm. The sender tries to transmit 20000 messages in succession to the corresponding receiver which counts the number of messages received.

For the success of this experiment, it was crucial that the MAC layer allows parallel transmission of multiple messages. Consequently, we adjusted the collision-preventing MAC layer delivered with TinyOS: Before sending a message, no check is performed if the medium

is free. Additionally, the initial random backoff before sending a message was removed. The output powers were fixed to 0 dBm for x_1 and -10 dBm for x_3 . We refer to this adjusted protocol as “SINR-MAC”.

Before presenting the measurement results, we calculate a theoretical lower bound for the time required to transmit the 20000 messages when assuming a graph-based communication model. A single packet containing 6 bytes of payload requires transmitting 23 bytes due to preamble, header, etc. The sensor sends with a data rate of 2.4 kbps and switching from RX to TX mode and back to RX mode requires about 0.5 ms, according to the CC1000 data sheet.³ Summed up, at least $(23 \text{ bytes}/2.4 \text{ kbps} + 0.5 \text{ ms}) * 40000 \text{ packets} \approx 403 \text{ s}$ are required when assuming a graph-based model. Note that this lower bound is very conservative, ignoring any software overheads.

We obtained a value of 603s for x_1 and 591s for x_3 using the standard TinyOS MAC layer protocol.⁴ In contrast, the “SINR-MAC” only required 267s (x_1) and 268s (x_3), respectively. This performance gain did not have negative effects on the reliability: With the standard MAC layer, x_2 received 19998 messages and x_4 received 18852 messages. The corresponding values using the “SINR-MAC” are 18668 for x_2 and 19916 for x_4 .

These results show that the examples analyzed in the previous section can be implemented in practice. On the one hand, the time used by the default MAC layer protocol exceeds the calculated lower bound by almost 50%. On the other hand, the “SINR-MAC” exploiting the interference phenomena of the SINR model performs significantly better than this limit. This highlights the inherent inability of graph-based models to represent important physical aspects that govern real sensor network. More importantly, however, this experiment shows that a protocol that is specifically tailored to the SINR model—in this case the adjusted “SINR-MAC” layer protocol—can outperform conventional, implicitly graph-based protocols by a factor of 2 or more.

3.2 Multiple Pairs of Nodes

Delighted by the results of the experiment above, the question arises if standard sensor nodes allow to use the medium threefoldly. The setup was analogous to the previous measurement with an additional sender and receiver pair, as shown in Figure 3. The output power of the senders was set to 0 dBm (x_1), -10 dBm (x_3), and -20 dBm (x_5).

The distances between the nodes were found by trial and error. During the search for promising distances,

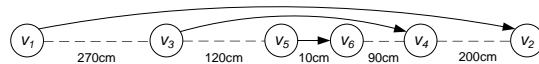


Figure 3: Three interleaved sender-receiver pairs.

we noticed that this setup is less failure tolerant than the first experiment with only two sender and receiver pairs: While moving a node a few centimeters to the left or to the right did not produce significant changes in the results, bigger changes led to complete failure of a receiver, i.e. it did no longer receive any messages destined for it. The reason is that in this experiment, each sender now has two competitors, whose interferences cumulate and reduce the region with sufficient SINR.

In this experiment, the same parameters as in the experiment above were measured, i.e. the time required by the three senders to completely send all 20000 messages and the number of successfully decoded messages.

	Time required	
	standard MAC	“SINR-MAC”
Node x_1	721 s	267 s
Node x_3	778 s	268 s
Node x_5	780 s	270 s

The number of successfully received messages at x_2 , x_4 , and x_6 using the standard MAC protocol was 19999, 18784, and 16519, respectively. For the “SINR-MAC”, the corresponding values were 19773, 18488, and 19498.

These measurements further emphasize the inability of graph-based approaches to model real sensor networks. The time required to send the 20000 messages is invariant even with three nodes sending messages while the standard MAC layer—as graph-based calculations would suggest—requires almost three times longer. Additionally, the number of collisions increased for the default MAC layer protocol resulting in a packet loss rate of 7.83% while the adjusted MAC layer shows a more or less invariant rate compared to the previous results.

Building systems with four or more senders transmitting messages in parallel becomes more and more impractical. On one side, this is because each additional sender increases the interference at the other senders. On the other side, the radio module only supports a limited interval of output powers. Our experiments have shown that four senders placed in a line are possible under perfect conditions. However, such systems tend to be very failure-prone in real environments. But different hardware platforms may produce different results.

4 APPLICATIONS & CHALLENGES

In the previous sections, we have seen how graph-based models inherently fail in capturing certain important physical phenomena. The fact that graph-based models do not properly describe all aspects of physical reality

³ChipCon AS, SmartRF CC1000 Datasheet (rev. 2.2), http://www.chipcon.com/files/CC1000_Data_Sheet_2.2.pdf

⁴All presented results are from one single run; however, we repeated all tests, and obtained similar results.

is of course neither new nor surprising (see for instance [2, 7, 10]). The more interesting question is whether the resulting inconsistencies are small enough to be justified by the gained simplicity of the model. Moreover, it is important to ask whether physical phenomena that cannot be modeled as graphs can be exploited for designing (and analytically proving!) algorithms and protocols that break the theoretical boundaries placed by graph-based models. In other words, we raise the question whether there are applications for wireless multi-hop networks in which provably efficient (possibly even theoretically optimal) graph-based algorithms perform inherently worse than algorithms that are designed to make use of SINR aspects. If there were no such examples, it would serve as a major justification for studying networks on the clean abstraction layer of graphs. If, however, there are examples in which the performance of theoretically optimal graph-based algorithms is surpassed by algorithms that explicitly take SINR into account, it would highlight the need for a more physical-level oriented design and analysis of network protocols.

One important application is *data gathering* with high throughput requirements in *heterogeneous wireless multi-hop networks*. Specifically, consider a heterogeneous network with potentially energy-restricted wireless nodes that gather data and locally distribute or forward this data for aggregation, and a few designated, more powerful nodes. Eventually, the data has to be sent to a base station, a task which is preferably done by the long-range nodes, instead of the regular sensor nodes. In any graph-based model, local communication among regular nodes and long-range communication among designated nodes must be coordinated (either in the time or frequency space, or by using spatial multiplexing). As in the four-node example of Figure 2, however, long-range and short-range communication can *co-exist*, that is, regular nodes can communicate with each other *while* long-range nodes send data to the base station. This could result not only in *higher throughput*, but also in a significantly *smaller coordination overhead* between different regions of the networks. Other applications could include improving the capacity in *wireless mesh networks* or even *cellular networks*.

In the remainder of this section, we want to theoretically study an application in which, even from an information-theoretic point of view, the theoretical limitations of graph-based models can be surpassed when explicitly using protocols designed for SINR environments.

Improving Channel-Throughput

Consider a *multi-hop channel* consisting of a chain of wireless nodes. The left-most node is the sender that wants to send data to the right-most node, its destination. Nodes being power constrained, the messages must

be forwarded in a multi-hop fashion from source to destination. The question is, at what *rate* R can data be transmitted in this model, that is, how much information can be successfully transmitted from source to destination in a certain time-interval.

In the formal model, the chain consists of n equidistantly placed nodes x_1, \dots, x_n , where x_1 and x_n are the source and destination, respectively. In the graph-based model, the *maximum transmission range* of any node is denoted by ℓ , $\ell < n$, i.e., a node x_i can send a message to $x_{i+\ell}$ in the absence of any interference.

We do not consider complex wireless signal propagation models because, interestingly, it suffices to study the basic *physical model* [8] in order to highlight the difference to graph-based models. Also, all our lower bounds hold even in very simplistic and optimistic graph models in which the interference range equals the transmission range, and in which time is divided into globally synchronized slots. Clearly, in more realistic graph models in which the interference range exceeds the transmission range, and in case of asynchrony, the achievable rates would be even worse.

We begin by showing that the naive idea to ship information from x_1 to x_n achieves only a very moderate rate. Consider the protocol in which every node transmits at power ℓ' , for some $1 \leq \ell' \leq \ell$. When having a message, a node x_i sends this message to $x_{i+\ell'}$ at the earliest opportunity. It can be seen in Figure 4(a) that x_1 can insert a message into the chain in time slot 1, but then has to wait for 2 consecutive time slots, before injecting its next message. The reason is that node $x_{\ell'+1}$ experiences interference during these two slots. As x_1 can thus insert a new message into the chain only once every three time slots, the achieved rate is $R = 1/3$.

Observation 1. *The rate achieved by the naive graph-based protocol of Figure 4(a) is $R = 1/3$.*

Clearly—even in graph-based models—much better protocols can be devised when using *power control*. Specifically, the rate can be improved by employing the forwarding scheme shown in Figure 4(b). Intuitively, if sending a message to its ℓ' -hop neighbor is impossible due to interference at the receiver, the message is forwarded to a closer neighbor where reception is possible. It can be shown that in this scheme, the channel allows the injection of 3 packets every 7 time slots.

Observation 2. *The rate achieved by the improved graph-based protocol of Figure 4(b) is $R = 3/7$.*

By using more complicated graph-based techniques, this rate may be improved further. However, the following theorem proves that even with the most sophisticated power control scheme and scheduling approach, the rate of $1/2$ can never be surpassed by a protocol that obeys the laws of a graph-based model.

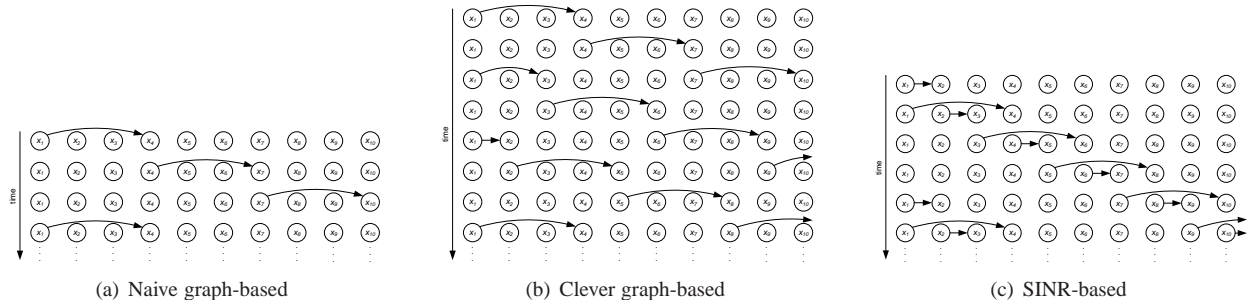


Figure 4: Figure 4(a) shows a naive, graph-based approach to send data from x_1 to x_n using $\ell' = 4$. A more sophisticated method to send messages from x_1 to x_n achieving a rate $R = 3/7$ is shown in Figure 4(b). The scheme in Figure 4(c) explicitly employs the SINR model to send messages from x_1 to x_n .

Theorem 3. *The maximum achievable rate R of graph-based scheduling protocol is*

$$R \leq \frac{1}{2} - \frac{1 - \frac{2\ell^2}{n}}{4(\ell + \frac{1}{2})}.$$

For $\ell < \sqrt{2n}/2$, this is strictly smaller than $1/2$.

Proof. Consider an arbitrary time slot. If in this time slot a node x_i transmits a message to node x_j that is h_i hops away ($j = i + h_i$), there cannot be any message sent to a node in the interval $x_{i-h_i}, \dots, x_i, \dots, x_{i+h_i}$. We call such nodes *blocked*, because no message can be sent to them. Let h_i be the number of hops node x_i transmits its message in this time slot, and let b_i denote the resulting number of blocked nodes. The following relation holds between h_i and b_i : $b_i = 2h_i + 1$ for $i \geq 2 + h_i$ and $b_i \geq h_i + 1$ for $i < 2 + h_i$.

Notice that there can be at most one successful sender x_i with $i < 2 + h_i$ in one time slot. Let S be the set of indices of successful senders in this time slot and let $P := \sum_{i \in S} h_i$ denote the amount of *progress* achieved by all nodes in the chain. The number of blocked nodes can be at most n , which implies that

$$i' + h' + 2 \cdot \sum_{i \in S \setminus \{x_{i'}\}} h_i + |S \setminus \{x_{i'}\}| \leq n, \quad (1)$$

where h' and i' are the hop-distance and the location of the left-most transmission, respectively. Because $i' \geq 1$ and $h' \leq \ell$, this can be rewritten as $\sum_{i \in S} h_i \leq \frac{1}{2}(n + \ell - |S|)$. On the other hand, it is clear that $\sum_{i \in S} h_i \leq \ell \cdot |S|$ holds because every node in S can at most send over ℓ hops. Hence, the progress P in every time slot is bounded by $P = \sum_{i \in S} h_i \leq \min\{\frac{1}{2}(n + \ell - |S|), \ell|S|\}$, which is maximized when $|S| = \frac{n + \ell}{2(\ell + 1/2)}$. Plugging in this value, the resulting progress is at most

$$P \leq \frac{n}{2} - \frac{n - 2\ell^2}{4(\ell + 1/2)}.$$

The theorem now follows because an algorithm with rate Λ must achieve a total progress of at least $tn \cdot \Lambda$ in t time slots, when $t \rightarrow \infty$. Because progress in each time slot is bounded by P , however, the achievable rate is at most P/n , which yields the claimed result of the theorem. \square

In view of this theorem, the question is whether $1/2$ is a fundamental barrier that cannot be surpassed by *any* protocol, or whether it is imposed solely by the underlying graph model. As it turns out, the latter is the case and depending on the values of α and β , the achievable rate can be at least $1/2$. In the scheme illustrated in Figure 4(c), for instance, x_1 first sends a packet to its one-hop neighbor. In the second iteration, this packet is forwarded one additional hop, to x_3 . Simultaneously—and this is where we abandon the graph-based model—the sender x_1 transmits a second packet to x_4 . As shown in Section 2, this is possible in the SINR model. Subsequently, these two messages are forwarded in the same manner in every time slot: the trailing message “hops” over the leading message, until they reach the destination. When doing the necessary calculations in the SINR model, it can be shown that for some α and β and appropriate power levels, this scheme can reach a rate of $R = 1/2$, because x_1 can inject a new packet in two time slots out of four. In fact, it may be the case that more sophisticated SINR-based schemes than the one shown in Figure 4(c) reach a rate strictly larger than $1/2$.

Observation 4. *For certain values of α and β , SINR-based scheduling protocols achieve a rate of $R \geq 1/2$.*

Notice that in its current form, the scheduling protocol of Figure 4(c) is valid only for large α and small β , and it may not be practically employable in certain settings for this reason. However, it serves to illustrate the potential gain in throughput by employing protocols and algorithms explicitly and making use of SINR phenomena. In particular, Observation 4 shows that by using the method of consecutively “overtaking” messages, the achievable rate can be $1/2$, whereas Theorem 3 proves that no graph-based protocol can achieve such a rate.

5 RELATED WORK

The discrepancy between graph-based models and physical SINR models has been recognized by researchers many years ago. For instance, the papers [2, 7] evaluate the performance of graph-based scheduling protocols in SINR environments by means of simulations and on the

assumption that nodes are distributed uniformly in the plane. Our work goes beyond these papers in the sense that we suggest to actually *design protocols* explicitly for SINR-based models, thus improving currently employed protocols. In fact, when it comes to scheduling, there already exist numerous algorithms in SINR environments, including for instance [4–6]. The authors of [4, 5] study the problem of finding a schedule and power control policy that minimizes the total average transmission power in the wireless multi-hop network. While these works provide important results, the proposed algorithms do either not yield efficient guarantees in arbitrary networks or are based on solutions to complex non-polynomial optimization problems. Computationally efficient solutions with provable guarantees that utilize SINR effects similar to the ones in this paper have only recently been studied for scheduling [10] and topology control [11].

In systems research, the general idea of exploiting “collision-but-not-failure” effects has been considered by Whitehouse et al., which makes explicit use of the *capture effect* [13]. Our work is different in that nodes actively select their power levels appropriate for creating desirable capture effects.

There has recently been a tremendous research effort towards increasing throughput in wireless networks, and some proposed strategies go beyond graph-based models [3, 9]. Biswas and Morris propose an improved routing and MAC layer protocol to enhance throughput in wireless networks [3]. Katti et al. propose an architecture for wireless mesh networks that disposes of the point-to-point abstraction of wireless channels and is based on the idea of *network coding* [9]. Since neither paper exploits SINR-effects at the *physical reception* of messages, they abandon the graph-based model on a different layer than we do. That is, applying ideas from network coding is completely *orthogonal* to our proposal, and hence can be applied in combination with SINR-based methods to further improve the results.

Another direction towards improving network capacity has been to use multiple or cognitive radios and allow communication on different frequencies, e.g. [1]. Again, these strategies are orthogonal to our work because SINR-effects can be exploited at each frequency individually.

6 CONCLUSIONS & OUTLOOK

Sections 3 and 4 have shown that protocols explicitly designed with SINR in mind can surpass the theoretically achievable performance of any graph-based protocol even in simple settings. The real challenge of course is to take these observations one step further, both theoretically and practically. From a *theoretical* point of view, it would be interesting to further characterize the gap between achievable rates in networks based on graph

models, the physical model, or even more realistic signal propagation models. While there is ample work dealing with exactly these kind of questions, we believe that there is still a lack of fundamental *algorithmic foundation* that allows to transform the theoretical insights into efficient and practical network protocols.

Even more important challenges, however, arise from *practical aspects* of our observations, i.e., turning them into practical network protocols. It would for instance be intriguing to study whether MAC layer protocols with higher throughput could be devised. Also, there is a large potential for improving the throughput of routing protocols or data gathering applications by incorporating our ideas. The ultimate challenge will be to circumventing the inevitably arising practical difficulties in order to tap the full potential of these technologies.

REFERENCES

- [1] P. Bahl, R. Chandra, and J. Dunagan. SSCH: Slotted Seeded Channel Hopping for Capacity Improvement in IEEE 802.11 Ad-Hoc Wireless Networks. In *Proc. of the 10th Int. Conference on Mobile Computing and Networking (MOBICOM)*, 2004.
- [2] A. Behzad and I. Rubin. On the Performance of Graph-based Scheduling Algorithms for Packet Radio Networks. In *Proc. of the IEEE GLOBECOM*, pages 3432–3436, 2003.
- [3] S. Biswas and R. Morris. ExOR: Opportunistic Multi-Hop Routing for Wireless Networks. In *Proc. of ACM SIGCOMM*, Philadelphia, Pennsylvania, USA, 2005.
- [4] R. L. Cruz and A. V. Santhanam. Optimal Routing, Link Scheduling and Power Control in Multi-hop Wireless Networks. In *Proc. 22th IEEE INFOCOM*, 2003.
- [5] T. ElBatt and A. Ephremides. Joint Scheduling and Power Control for Wireless Ad-hoc Networks. In *Proceedings of the 21th IEEE INFOCOM*, 2002.
- [6] J. Grönkvist. Assignment Strategies for Spatial Reuse TDMA. *Licentiate thesis, Royal Institute of Technology, Stockholm, Sweden*, 2002.
- [7] J. Grönkvist and A. Hansson. Comparison Between Graph-Based and Interference-Based STDMA Scheduling. In *Proc. 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MOBIHOC)*, pages 255–258, 2001.
- [8] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.
- [9] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft. XOR in The Air: Practical Wireless Network Coding. In *Proc. of ACM SIGCOMM*, Pisa, Italy, 2006.
- [10] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. of the 25th IEEE INFOCOM*, 2006.
- [11] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology Control meets SINR: The Scheduling Complexity of Arbitrary Topologies. In *Proc. of the 7th ACM Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2006.
- [12] D. Son, B. Krishnamachari, and J. Heidemann. Experimental Analysis of Concurrent Packet Transmissions in Low-Power Wireless Networks. Technical report, Viterbi School of Engineering, University of Southern California, 2005.
- [13] K. Whitehouse, K. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the Capture Effect for Collision Detection and Recovery. In *Proc. of the 2nd IEEE Workshop on Embedded Networked Sensors (EmNetS-II)*, pages 45–52, 2005.