Distributed Computing Column 38

Models for Algorithm Design in Wireless Networks

Idit Keidar
Dept. of Electrical Engineering, Technion
Haifa, 32000, Israel
idish@ee.technion.ac.il

The prosperity of research on wireless communication has not skipped the distributed computing community. Wireless networks provide unique and challenging platforms for distributed computation, inspiring researchers to develop many distributed algorithms for such environments. Any algorithmic work targeting wireless networks must begin by defining an appropriate model. The first research results in this vein employed simplified models, like the *Unit Disk Graph (UDG)*, which facilitated algorithm design and were readily amenable to analysis. Recently, models that more accurately capture the physical nature of wireless networks were developed; such models are nowadays being gradually adopted.

This column's contribution, by Zvi (Zvika) Lotker and David Peleg focuses on the promising *signal-to-interference & noise ratio (SINR)* model, and studies it from an algorithmic perspective. It surveys the model's structural properties as well as algorithms designed for this model. Along the way, some of the similarities and differences between the SINR model and the UDG model are highlighted. Many thanks to Zvika and David for their contribution!

Call for contributions: I welcome suggestions for material to include in this column, including news, reviews, opinions, open problems, tutorials and surveys, either exposing the community to new and interesting topics, or providing new insight on well-studied topics by organizing them in new ways.

Structure and Algorithms in the SINR Wireless Model

Zvi Lotker[*]
Dept. of Communication Systems Engineering
Ben Gurion University
Beer-Sheva, Israel
zvilo@cse.bgu.ac.il

David Peleg
Dept. of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot, Israel
david.peleg@weizmann.ac.il

Abstract

The *signal-to-interference & noise ratio (SINR)* model is one of the most commonly studied physical (or fading channel) models for wireless networks. We survey some recent studies aiming at achieving a better understanding of the SINR model and its structural properties and developing efficient design algorithms and communication protocols for it.

Physical models and SINR

Traditional (wired, point-to-point) communication networks can be described satisfactorily using a graph representation. A station s is able to transmit a message to another station s' if and only if there is a wire connecting the two stations. This condition is independent of the locations, connections and activities of the two stations or of other nearby stations[1]. Accurately representing a wireless network is considerably harder, since it is nontrivial to decide whether a transmission by a station s is successfully received by another station s'; this may depend on the positioning and activities of s and s', and on other nearby stations, whose activities might interfere with the transmission and prevent its reception This means that a transmission from s may reach s' in some settings but fail to reach it under other settings. Moreover, the question of successful reception is more complex, since connections can be of varying quality and capacity.

In fact, there are many other relevant factors, such as the presence of physical obstacles, the directions of the antennae at s and s', the weather, and more. Obtaining an accurate solution taking all of those factors into account involves solving the corresponding Maxwell equations. Since this is usually far too complicated, the common practice is to resort to approaches based on approximation models. For instance, one way to predict the behavior of wireless systems in an urban environment is using a ray tracing model, based on the assumption that radio waves behave according to geometric optics where walls are modeled as reflective mirrors. The uncertainly involved in the dynamics of a radio channel can be modeled using a Markov process. For more information we refer the interested reader to Chapters 2 and 3 of [9].

For the purposes of the current discussion, we follow the approach of ignoring those complicating factors, and assuming a relatively clean abstract setting where the only players are the transmitting and listening

[*]Partially supported by a gift from Cisco research center and by the Israel Science Foundation, grant 894/09.
[1]excepting certain types of wired local area networks.

stations, and the antennae are omni-directional. In this setting, the rules governing the reception quality of wireless transmissions can be described schematically by *physical* or *fading channel* models. Among those, one of the most commonly studied is the *signal-to-interference & noise ratio (SINR)* model. In this model, the energy of a signal fades with the distance to the power of the path-loss parameter α. When a station s transmits, the message is successfully received by a listening station s' if and only if the strength of the signal received by s', divided by the strength of the interferences from other simultaneous transmissions (plus the background noise N), exceeds some threshold β. Hence for a collection $S = \{s_1, \ldots, s_n\}$ of simultaneously transmitting stations in the plane, it is possible to identify with each station s_i a *reception region* $\mathcal{H}(s_i)$ around it, consisting of the points where the transmission of s_i is received correctly. More precisely, denote by $\mathrm{dist}(p, q)$ the Euclidean distance between p and q, and assume that each station s_i transmits with power E_i. At an arbitrary point p, the transmission of station s_i is correctly received if

$$\frac{E_i \cdot \mathrm{dist}(p, s_i)^{-\alpha}}{N + \sum_{j \neq i} E_j \cdot \mathrm{dist}(p, s_j)^{-\alpha}} \geq \beta .$$

This formula represents a rather general model concerning the allowed transmission power, referred to as the *power control* model, in which each station can control the power with which it transmits. A simpler (and weaker) model is the *uniform* wireless network model, which assumes that all transmissions use the same amount of energy, i.e., $E_i = 1$ for every i.

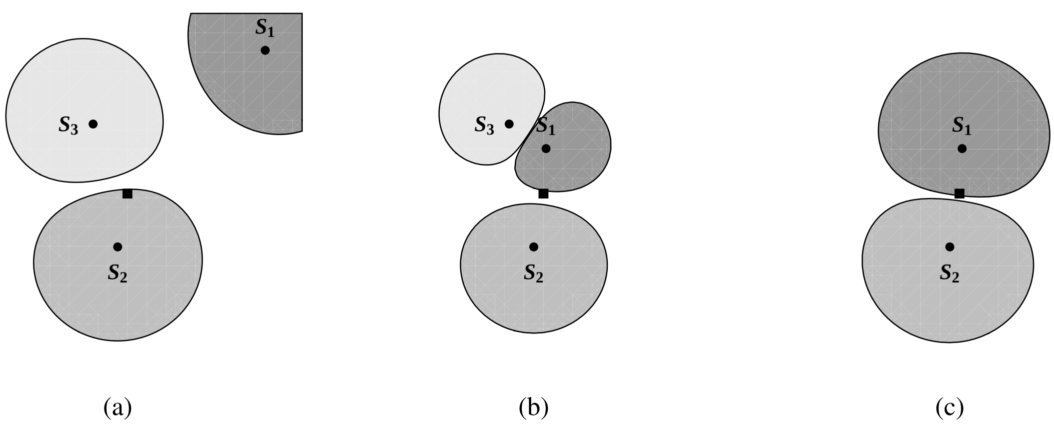

Figure 1: An SINR network with three simultaneously transmitting stations s_1, s_2, s_3 and one receiver, marked by the solid black square. (a) The receiver can hear s_2. (b) Station s_1 is relocated and now the receiver cannot hear any of the transmissions. (c) If, at the same locations as in (b), s_3 remains silent while s_1 and s_2 transmit, then the receiver can hear s_1.

The dependence of reception on the stations' locations and activities is illustrated in Figures 1 and 2. Fig. 1 depicts a uniform wireless network consisting of three stations s_1, s_2, s_3 and their reception regions under various settings. Figure 2 illustrates more involved configurations that may arise in the non-uniform case or in the uniform case with $\beta < 1$.

Graph-based models and UDG

A fair amount of research exists on the SINR model and other variants of the physical model, yet as far as algorithms are concerned, progress has been rather slow. Among the reasons for this are the fact that in physical models it is nontrivial to decide some basic questions on a given setting, and it is much harder to develop

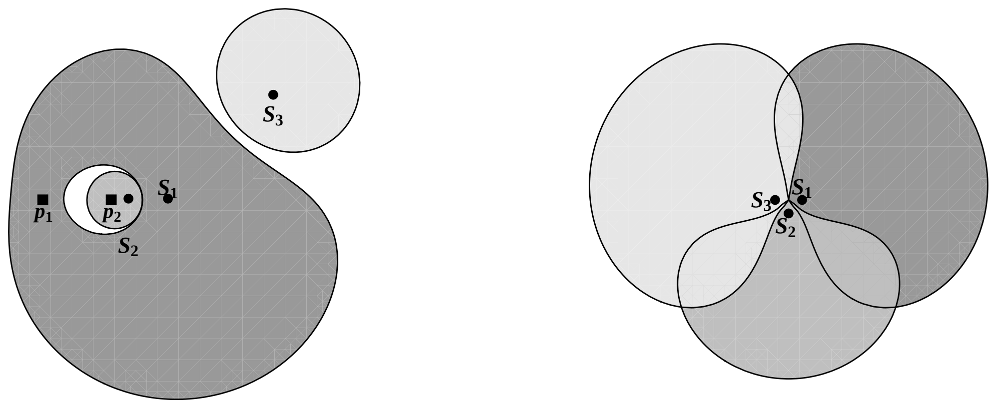

(a) (b)

Figure 2: Reception regions in the SINR model. (a) A non-uniform network where station s_1 transmits with a higher power level than the other two stations. The white regions (including the crescent-like internal one) represent “no-reception” areas. (b) A uniform network with path-loss parameter $\alpha = 2$, reception threshold $\beta = 0.3$ and background noise N = 0.05.

and analyze communication protocols and network design algorithms. Consequently, previous research of wireless multi-hop networks, including the study of issues such as transmission scheduling, frequency allocation, topology control, connectivity maintenance, routing, and related design and communication tasks, has relied on simplified *graph-based* models rather than on physical models.

Graph-based models represent the network by a graph $G = (S, E)$ such that a station s successfully receives a message transmitted by a station s' if and only if s and s' are neighbors in G and s does not have a concurrently transmitting neighbor in G. In particular, the model of choice for many protocol designers is the *unit disk graph (UDG)* model [14]. In the general *disk graph* model (a.k.a. the *protocol* model), the stations are represented as points in the Euclidean plane, and the transmission of a station s_i is received (in the absence of collisions with other simultaneous transmissions) by every other listening station v located within a disk centered around s_i, whose radius r_i depends on the energy E_i with which s_i transmits. Hence the reception relationships in a wireless network can be modeled by a *disk graph*, namely, a directed graph whose vertices correspond to the stations, with a directed edge leading from s_i to s_j if the two stations are at distance at most r_i from each other. Most of the literature on graph-based models for wireless networks concentrates on the *uniform* variant of this model, known as the UDG model, where all disk radii are assumed to be 1, hence the network is represented by an (undirected) unit disk graph [5], with an edge connecting any two stations whose corresponding points are at distance at most one from each other.

Comparisons between the SINR and UDG models

Let us illustrate some of the differences between the UDG and SINR models, with respect to the way they deal with interference. The UDG model provides a rather rigid representation for the reception region of a station s, taking it to be a unit disk around s. Hence, considering for example the configuration depicted in Fig. 3(a), we observe that a listening station v located at the solid black square is adjacent to the transmitting station s_1, but not to any of the other simultaneously transmitting stations, hence the UDG model predicts that v will receive the message in this configuration. In reality, the presence of several nodes slightly outside the unit disk around the receiver station v might still generate enough cumulative interference to prevent

v from successfully receiving a message from s_1, hence the UDG model might yield a "false positive" indication of reception. Indeed, the prediction of the SINR model in this same configuration, depicted in Fig. 3(b), is that the reception regions will be smaller than unit disks, and subsequently, v will not hear any of the transmitting stations.

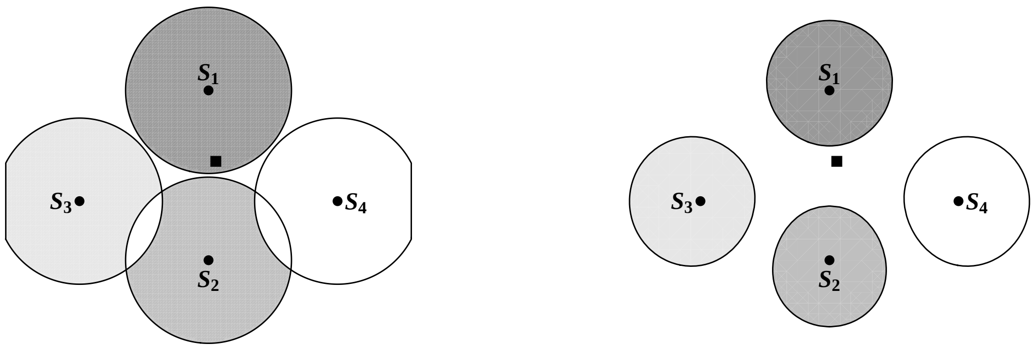

(a) (b)

Figure 3: (a) In the UDG model, p can hear s_1. (b) In the SINR model, the cumulative interference of stations s_2, s_3, s_4 prevents p from hearing s_1.

Conversely, a simultaneous transmission by two or more neighbors should not always end in collision and loss of the message; in reality this depends on other factors, such as the relative distances and the relative strength of the transmissions. Thought provoking experimental results presented in [20] show that even basic wireless stations can achieve communication patterns that are impossible in disk graph-based models. One such example is the pattern depicted in Figure 2(a), restricted to the stations s_1 and s_2 and the points p_1 and p_2. This configuration extends to two dimensions a 1-dimensional example of [20], and illustrates a possible setting in the SINR model in which p_1 can hear s_1 and p_2 can hear s_2 simultaneously. In contrast, it can be shown that in the disk graph model there can be no assignment of transmission powers that will achieve this effect. In the same spirit, certain situations are shown in [17] in which it is possible to apply routing / transport schemes that may break the theoretical throughput limits of any protocol that obeys the laws of a graph-based model.

An additional such scenario is illustrated in Figure 4, which compares the reception regions of the UDG and SINR models with four transmitting stations s_1, s_2, s_3, s_4 and one receiver p (represented as a solid black square). Suppose that initially, only station s_1 transmits, and all others remain silent. In this case, both models provide the same picture, in which p can hear s_1. Figure 4 illustrates the process of gradually adding s_2, s_3 and s_4 to the transmitting set.

Bridging between the models

Graph-based models are attractive for higher-layer protocol design, as they conveniently abstract away some of the complications of the wireless model. For example, issues of topology control, scheduling and allocation are handled more directly, since notions such as adjacency and overlap are easier to define and test, in turn making it simpler to employ also some useful derived concepts such as domination, independence, clusters, and so on. On the down side, it should be realized that graph-based models ignore or oversimplify a number of important physical aspects of real wireless networks, such as the laws of interference. This lim-

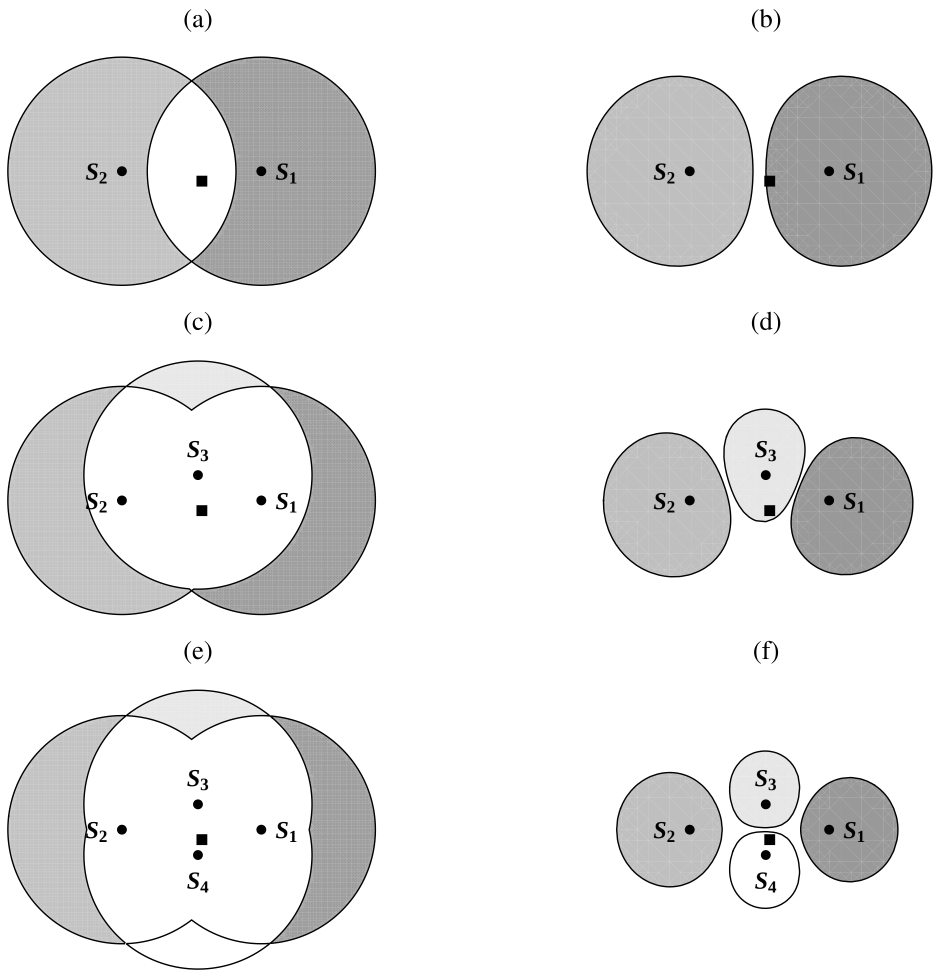

Figure 4: Reception regions in the UDG model (left) and SINR model (right). (a)-(b): When both s_1 and s_2 transmit simultaneously, p hears neither station in the UDG model, but it does hear s_1 in the SINR model, hence the UDG model yields a "false negative" indication. (c)-(d): When s_3 joins the transmitting stations, p still cannot hear any station in the UDG model, but now it can hear station s_3 in the SINR model. (e)-(f): Station s_4 starts to transmit as well.

its the applicability or accuracy of many of the algorithms presented in the literature for wireless networks under the UDG model.

In summary, while the existing body of literature on models and algorithms for wireless networks represents a significant base containing a rich collection of ideas, paradigms, tools and techniques, the limitations described above leave us in the unfortunate situation where the more practical graph-based models (such as the UDG model) are not sufficiently accurate, and the more accurate physical models (such as the SINR model) are not sufficiently practical or well-understood. This state of affairs has led to recent interest in the study of some of the structural and algorithmic aspects of the SINR model, as well as of the relationships between physical and graph-based models, in the hope of bridging the gap between these models. Ultimately, the motivation is to identify more suitable "hybrid" models, attaining some desirable tradeoff points between practicality and accuracy, and subsequently develop better and more efficient algorithmics for fundamental problems and tasks in wireless networks.

More elaborate graph-based models may employ two separate graphs, a *connectivity graph* $G_C = (S, E_C)$ and an *interference graph* $G_I = (S, E_I)$, such that a station s successfully receives a message transmitted by a station s' if and only if s and s' are neighbors in the connectivity graph G_C and s does not have a concurrently transmitting neighbor in the interference graph G_I. Protocol designers sometimes consider special cases of this general model. For example, it may be assumed that G_I is G_C augmented with all edges between 2-hop neighbors in G_C. An alternative variant of the UDG model handling transmissions and interference separately, named the ***Quasi Unit Disk Graph (Q-UDG)*** model, is introduced in [15]. In this model, two concentric circles are associated with each station, the smaller representing its reception region and the larger representing its area of interference. Another interference model, also based on the UDG model, is proposed in [24].

One way to bridge the gap between the two models is using emulation. Lebhar et al. [16] consider the case of $\alpha > 2$ and nodes that are deployed uniformly at random in a given area. For this setting, they show how a UDG protocol can be emulated when the network operates under the SINR model, with an emulation cost factor of $O(\log^3 n)$. Note, however, that this does not resolve all questions related to the SINR model, for three reasons. First, the paper only talks about randomly and uniformly deployed transmitters. Second, there is an overhead of $O(\log^3 n)$ in time. Finally, the algorithm works only for $\alpha > 2$.

A natural question concerns the difference between the power control model and the uniform model. It is possible to prove that if some resource (e.g., the used energy or the general area in which the network resides) is bounded, then the ratio between the two models is proportional to $\log B$, where B is the bound on the resource; see [1, 3, 13, 22].

Capacity and scheduling in the SINR model

Some recent studies aim at achieving a better understanding of the SINR model. In particular, in their seminal work [12], Gupta and Kumar analyze the capacity of wireless networks in the physical and protocol models. Specifically, they consider a setting where n transmitter-receiver pairs are located in the unit square. In this setting, they show that the capacity of the transmitters under any deployment where the average distance is constant is the same in both models. In contrast, Moscibroda [17] analyzes the worst-case capacity of wireless networks in arbitrary deployments (specifically, allowing arbitrary average distances), and establishes an exponential gap between the SINR model and the UDG model, by showing that there are transmitter deployments for which the capacity in the SINR model is exponentially greater than in the UDG model. (Those deployments involve some very short and some very long links, and the average distance when all stations are located in the unit square is $O(1/n)$.) In both studies, the deployments are not required to ensure connectivity. (Actually, it may be impossible to ensure connectivity if one considers only a single radio frequency and a single time step.)

Transmission scheduling (so as to prevent collisions) is a central issue in wireless communication theory. Given a set of transmitter-receiver pairs, the goal is to select a schedule for the transmissions, so that all the receivers correctly decode their messages. Several variants of the scheduling problem were studied in the literature. Some of those variants assume a given power assignment for the transmitters (which in some settings is required to adhere to a certain policy, such as uniformity or distance-proportionality, and in other cases is assumed to be arbitrary). Other variants require the algorithm to decide also on the power assignment for the transmitters (again, possibly adhering to some specific policy).

Schedules for basic network structures, namely, strongly connected networks, are studied in [19, 21]. It is shown that allowing arbitrary adjustments to the transmission power gives an exponential advantage over the uniform or linear power assignment schemes. (The latter require a transmitter s sending a message to a receiver r to use power proportional to $d(s,r)^\alpha$, the distance to power α.) This gives an interesting complement to the capacity bounds of [12, 17] discussed above.

A measure called *disturbance*, capturing the intrinsic difficulty of finding a short schedule for a problem instance, is defined in [18], along with an algorithm that achieves provably efficient performance (in terms of schedule length) in any network and request setting that exhibits low disturbance.

For the special case of many-to-one communication with data aggregation in relaying nodes (e.g., summing up some counters maintained at the nodes), a scaling law describing the achievable rate in arbitrarily deployed sensor networks is derived in [17]. It is shown that for a large number of aggregation functions, a sustainable rate of $1/\log^2 n$ can be achieved.

Goussevskaia [10] studies some additional algorithmic aspects of scheduling in the SINR model, and in particular presents an $O(\log(n))$ approximation algorithm for scheduling wireless requests. On the negative side, some impossibility results are proved in [11], and the NP-hardness of the scheduling problem in the SINR model is established in [10, 23].

In summary, the literature reviewed above presents several examples where the SINR model allows better performance than the UDG model, typically through solutions involving exponential gaps in transmission power or in distances between different transmitter-receiver pairs. We next discuss some further results derived specifically in the SINR model.

Several recent papers consider oblivious transmission scheduling schemes. In oblivious scheduling, the transmission power is entirely determined by the distance between the transmitter and the intended receiver. (In particular, $E(r) = c$ for constant c is a function determined by the distances, so uniform networks employ an oblivious scheme.) In [7], Fanghänel et al. prove that there is an exponential gap between the directed and bidirectional versions of the scheduling problem. Specifically, they show that in the case where each pair of points consists of a dedicated transmitting device and a dedicated receiving device for any oblivious power assignment scheme, there exists an instance with n directed communication requests that requires $\Omega(n)$ time steps. In contrast, when the communication is symmetric and both endpoints use the same transmission power, it is shown that there exists a universally good oblivious power assignment. This means that the capacity increases when we allow transmission power to vary between the different sides of a communication channel.

The problem of optimizing a single scheduling step was studied in [1, 6]. Given a collection of sender-receiver pairs (s_i, t_i) in the plane, the goal is to assign each transmitter s_i a power (possibly 0) so as to maximize the number of successful pairs. Let $\Delta = \max_i d(s_i,t_i)/\min_i d(s_i,t_i)$, i.e., the "aspect ratio" restricted to the sender-receiver pairs. NP-hardness of the problem is established in [1], as well as an $O(\log \Delta)$-approximation algorithm and a proof that in an appropriately defined game every Nash equilibrium has an expected number of successful connections that is within $O(\Delta^{2\alpha})$ of optimal. The main result of [6] is that if all transmitters use no-regret algorithms to play the game defined in [1], then the average number of successful connections (over a sufficient number of rounds) will be within $O(\Delta^{2\alpha})$ of optimal. This is the first distributed algorithm (in some sense) for this problem.

In [8], Fanghänel et al. introduce an instance-based measure I of interference, and prove for general power assignments in the plane a lower bound of $\Omega(\frac{I}{\log \Delta})$ steps for $\alpha > 2$. They also show that when restricted to linear power assignments, the bound can be improved to $\Omega(I)$; in this case they also present an efficient algorithm for computing an $O(I \log n)$ step schedule for linear power assignments. They also extend these results towards multi-hop scheduling and routing. In the multi-hop scheduling problem, the request consists of a sequence of pairs, referred to as *paths*, rather than a single pair of nodes.

In the context of routing, the problem of constructing minimum delay end-to-end schedules for a given set of routing requests is studied in [4]. In this problem, each node is assigned a distinct power level, the paths for all requests are determined, and all message transmissions are scheduled to guarantee successful reception in the SINR model. A polynomial-time algorithm with provable worst-case performance for the problem in this setting is presented in [4].

Another line of research, in which known results from the UDG model are analyzed under the SINR model, is explored in [21], which studies the problem of topology control. The stations are assumed to be embedded in the Euclidean plane, and are required to simulate a given underlying graph topology. For each station v, let r_v denote the length of the longest outgoing edge of v, and consider a ball of radius r_v centered at v. Denote the number of nodes contained in this ball by $I_{in}(v)$, and let $I_{in} = \max_{v \in V} I_{in}(v)$. The main result of [21] is that there exists a schedule that allows the communication to flow on the original topology and completes in time $O(I_{in} \log^2(n))$.

SINR diagrams

An issue of increasing significance is that of obtaining a "reception map" depicting the behavior of a multi-station wireless network. Such a reception map characterizes the reception regions of the stations, namely, partitions the plane into n reception regions $\mathcal{H}(s_i)$, $1 \le i \le n$, and a region $\mathcal{H}_\emptyset$ where no station can be heard. (Each of these $n+1$ regions may possibly be composed of several disconnected sub-regions). The map may change dynamically with time, as the stations may choose to transmit or keep silent, or even change their location from time to time.

In the UDG model, the basic (static) underlying *UDG diagram* of a given wireless system can be constructed in linear time in a straightforward manner. Moreover, for a given time step in which only some of the stations transmit simultaneously, and for a given point p, it is simple to decide which station (if any) is heard at p. The task becomes more challenging in the SINR model, where it is not completely clear what shapes the reception regions may take; in fact, it is not easy to construct an *SINR diagram* even in a static setting. Figures 1 and 2 illustrate the complexity of understanding SINR diagrams.

SINR diagrams appear to be fundamental to understanding the dynamics of wireless networks, and play a key role in the development of suitable algorithmics for such networks, analogous perhaps to the role played by Voronoi diagrams in the study of proximity queries and related issues in computational geometry. In fact, the analogy between SINR diagrams and Voronoi diagrams runs deeper; it can be shown that the SINR diagram of a set of points S converges to the Voronoi diagram of S when the path loss parameter α tends to infinity and $N = 0$. The case where $N > 0$ is a little more complicated and has a nice interpolation. In this case, one may compute the Voronoi cell of each station and intersect the cell with the a unit disk center at the transmitter. The resulting structure consists of regions known as alpha complexes / shapes. One use of the SINR diagram is to bridge between the Voronoi diagram and the structure of alpha complexes, as can be seen in Figure 5.

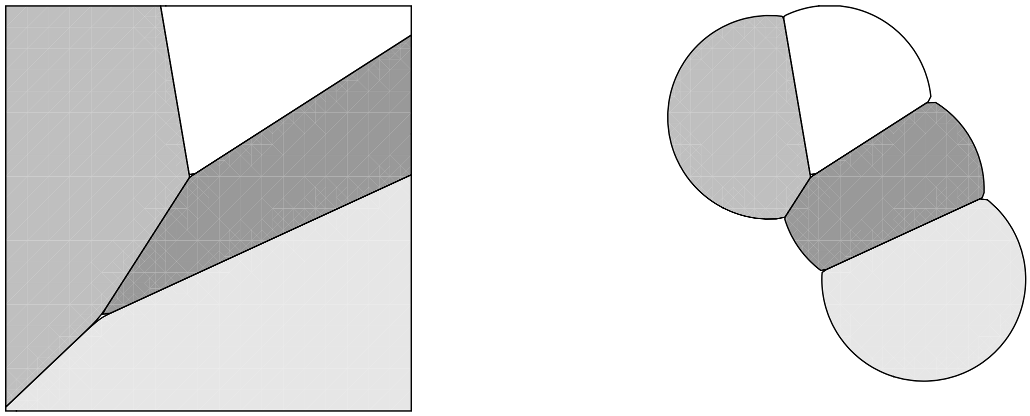

Voronoi (a) alpha (b)

Figure 5: As $\alpha \to \infty$, the SINR diagram converges (a) to the Voronoi diagram, when the background noise $N = 0$; (b) to the structure of alpha complexes, when $N > 0$.

SINR diagrams were introduced in [2], where it is shown that in uniform networks with $\alpha = 2$ and $\beta \geq 1$, every reception region in the SINR diagram is convex and *fat*[2]. This "niceness" of the reception regions may potentially make it easier to handle them by searching, sampling and location algorithms. In a certain sense, this result lends support to the model of *Quasi Unit Disk Graphs* suggested by Kuhn et al. in [15]. In contrast, when $0 < \beta < 1$, the reception regions may be nonconvex and even overlapping, as can be seen in figure 2(b).

The characterization of the reception regions as convex and fat is subsequently exploited in [2] for developing an efficient structure for answering point location queries, by preprocessing the station locations (in polynomial time) and providing each station with an additional data structure of size $O(1/\epsilon)$ that approximates the reception region boundaries, so that for every point outside the ϵ vicinity of the boundary (for performance parameter $\epsilon > 0$) it is possible to answer a point location query in $\log(1/\epsilon)$ time. Specifically, given a collection of simultaneously transmitting stations S, a station $s \in S$, and a point p, the data structure responds with one of three possible answers: (1) s is heard at p; (2) s is not heard at p; or (3) cannot say for sure. Letting $\mathcal{H}$ be the reception region of s and $\mathcal{F}$ be the region for which answer (3) is returned, the algorithm guarantees that $Area(\mathcal{F}) \leq \epsilon \cdot Area(\mathcal{H})$.

Conclusions

In this survey we discussed some recent studies of the SINR model, its structural properties and algorithms for it. Special attention was given to studies focusing on identifying the similarities and differences between the SINR model and the graph-based UDG model, and drawing analogies and algorithmic ideas in both directions.

While this area has attracted considerable attention in the last decade or so, much remains to be done, especially when it comes to algorithms for the SINR model. It is hoped that improving our understanding of the SINR model may lead to more efficient algorithms for a variety of design and communication problems.

In particular, the study of SINR reception diagrams and their structure and algorithms is still in its infancy. Hence a challenging direction involves broadening our scope of understanding concerning SINR

[2] i.e., the radius ratio between its largest enclosed circle and smallest enclosing circle is $O(1)$.

diagrams, allowing us to generate, interpret and utilize such diagrams for non-uniform power models. This may lead to the development of more accurate, yet practically usable, abstract models for wireless communication, and possibly to bridging the gap between physical and graph-based models.

References

[1] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th Conf. of IEEE Computer and Communications Societies (INFOCOM)*, 2009.

[2] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: towards algorithmically usable SINR models of wireless networks. In *Proc. 28th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 200–209, 2009.

[3] C. Avin, Z. Lotker, and Y.-A. Pignolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *ESA*, pages 373–384, 2009.

[4] D. Chafekar, V.S.A. Kumar, M. Marathe, S. Pathasarathy, and A. Srinivasan. Cross-layer latency minimization in wireless networks with SINR constraints. In *Proc. 8th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2007.

[5] B.N. Clark, C.J. Colbourn, and D.S. Johnson. Unit disk graphs. *Discrete Math.*, 86:165–177, 1990.

[6] M. Dinitz. Distributed algorithms for approximating wireless network capacity. In *Proc. 29th Conf. of IEEE Computer and Communications Societies (INFOCOM)*, 2010.

[7] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *Proc. 28th ACM Symp. on Principles of distributed computing (PODC)*, pages 220–229, 2009.

[8] A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *Proc. 36th Int. Colloq. on Automata, Languages and Programming (ICALP)*, pages 447–458. SV, 2009.

[9] A. Goldsmith. *Wireless communications.* Cambridge University Press, 2005.

[10] O. Goussevskaia. *Computational Complexity and Scheduling Algorithms for Wireless Networks*. PhD thesis, ETH Zurich, 2009.

[11] O. Goussevskaia, Y.-A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Proc. 8th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 100–109, 2007.

[12] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Trans. Information Theory*, 46:388–404, 2000.

[13] Magnús M. Halldórsson. Wireless scheduling with power control. In *Proc. 17th European Symp. on Algorithms (ESA)*, pages 361–372, 2009.

[14] M.L. Huson and A. Sen. Broadcast scheduling algorithms for radio networks. In *IEEE Military Communications Conf. (MILCOM)*, pages 647–651, 1995.

[15] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-Hoc Networks Beyond Unit Disk Graphs. In *1st ACM Workshop on Foundations of Mobile Computing (DIALM-POMC)*, 2003.

[16] E. Lebhar and Z. Lotker. Unit disk graph and physical interference model: Putting pieces together. In *23rd IEEE Int. Symp. on Parallel and Distributed Processing (IPDPS)*, 2009.

[17] T. Moscibroda. The worst-case capacity of wireless sensor networks. In *Proc. 6th Int. Conf. on Information Processing in Sensor Networks (IPSN)*, pages 1–10, 2007.

[18] T. Moscibroda, Y.A. Oswald, and R. Wattenhofer. How optimal are wireless scheduling protocols? In *Proc. 26th Conf. of IEEE Computer and Communications Societies (INFOCOM)*, 2007.

[19] T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th Conf. of IEEE Computer and Communications Societies (INFOCOM)*, 2006.

[20] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *Proc. 5th Workshop on Hot Topics in Networks (Hotnets)*, 2006.

[21] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology control meets SINR: the scheduling complexity of arbitrary topologies. In *Proc. 7th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 310–321, 2006.

[22] M. Parter and D. Peleg. Bounding the power-to-spread ratio in SINR wireless networks. Unpublished manuscript, 2010.

[23] Y.-A. Pignolet. *Algorithmic Challenges in Wireless Networks Interference, Energy and Incentives*. PhD thesis, ETH Zurich, 2009.

[24] P. von Rickenbach, S. Schmid, R. Wattenhofer, and A. Zollinger. A robust interference model for wireless ad-hoc networks. In *Proc. 19th Int. Parallel and Distributed Processing Symp. (IPDPS)*, 2005.