A Primer on Spatial Modeling and Analysis in Wireless Networks

Jeffrey G. Andrews and Radha Krishna Ganti, The University of Texas at Austin
Martin Haenggi, University of Notre Dame
Nihar Jindal, University of Minnesota
Steven Weber, Drexel University

ABSTRACT

The performance of wireless networks depends critically on their spatial configuration, because received signal power and interference depend critically on the distances between numerous transmitters and receivers. This is particularly true in emerging network paradigms that may include femtocells, hotspots, relays, white space harvesters, and meshing approaches, which are often overlaid with traditional cellular networks. These heterogeneous approaches to providing high-capacity network access are characterized by randomly located nodes, irregularly deployed infrastructure, and uncertain spatial configurations due to factors like mobility and unplanned user-installed access points. This major shift is just beginning, and it requires new design approaches that are robust to spatial randomness, just as wireless links have long been designed to be robust to fading. The objective of this article is to illustrate the power of spatial models and analytical techniques in the design of wireless networks, and to provide an entry-level tutorial.

THE IMPORTANCE OF SPACE IN WIRELESS NETWORKS

The importance of transmit-receive distance in wireless communication has long been known. For example, the range of a wireless link has been considered important dating to the time of Marconi, and such thinking has evolved into link-budget analyses that are buttressed by fading margins, resulting in well-understood rate vs. range trade-offs [1]. While such an approach is appropriate for point-to-point wireless links, it is insufficient for networks due to the critical role of interference. Within the literature on wireless network design, the desire for tractability has led to over-simplified models, such as assuming all interfering nodes contribute equally to the aggregate interference or disc-like models where interferer impact is binary, i.e., interference has either no impact or it leads to complete packet loss. Such models fail to correctly capture the nature of wireless propagation or the importance of the interferers' locations, since average received signal strength falls off *continuously* with distance.

The challenge of spatial modeling can be illustrated by comparing the spatial resource to time/frequency resources. Wireless transmissions need to be separated in time, frequency, and/or space to avoid excessive interference. Space is by far the most challenging resource to use efficiently, for two reasons:

- In space, transmitters and receivers are not collocated
- Power from undesired transmitters leaks in space over relatively large distances

In contrast, when using time or frequency division, transmitters and receivers are collocated (in time/frequency), and the spilling can be minimized. This is illustrated in Fig. 1. In time, when turning off a transmitter, radiated power is driven to zero almost immediately. In frequency, waveforms are designed such that the power fall off is 100 dB or more per decade. In contrast, the falloff in space is only about 20–40 dB/decade. For example, Wi-Fi devices are required to suppress their transmissions in frequency by 20 dB over just 2 MHz, which is only 10 percent of their bandwidth and 0.01 percent of a decade (which would run to 24 GHz). In short, interference in time and frequency can be engineered, whereas interference in space is held hostage by Maxwell's equations and there are very few practical options available to reduce interference to neighboring receivers short of reducing the transmit power; which would equally affect the strength of the desired signal and thus not increase the signal-to-interference ratio.

RANDOM SPATIAL MODELS

Because spatial configurations may vary widely over an enormous (often infinite) number of possibilities, one cannot design most systems for each specific configuration but must instead consider a *statistical* spatial model for the node locations. The usefulness of recent innovations such as wireless network coding and interference alignment depends critically on the relative positions of transmitters and receivers: but just how likely are configurations where these techniques are effective? An accurate performance assessment is only possible with an accurate probabilistic model for those positions. Just as one uses a fading or shadowing distribution to model the variety of possible propagation environments in a wireless link, one should use a statistical distribution to model the variety of possible network topologies. In spite of this necessity, accounting for the distribution on locations has been somewhat slow to migrate into our understanding of wireless networks.

Femto, Cognitive, and Mesh Networks

Many of the most pressing questions about wireless networks are fundamentally spatial in nature, for example:

- What effect will the largely unplanned deployment of hundreds of millions of femtocells, relays, and other hotspots in the next decade have on other users of that spectrum?
- When and how can cognitive radios be used in white space without affecting the performance of incumbent users, who are typically distributed in space with unknown locations?
- How should wireless multihop networks be designed, given that users are scattered in space and may even be mobile? What are good physical layer, medium access control (MAC), and routing protocols for ad hoc and mesh networks?

Fortunately, there is a sophisticated mathematical toolset that can be applied to help answer such questions. The field of stochastic geometry (i.e., the statistical modeling of spatial relationships) has been extensively developed over the last two decades and continues to mature with a flurry of recent results extending up to the present. Although most applications of stochastic geometry have thus far focused on purely wireless (ad hoc) networks, these tools are equally useful and necessary for other emerging wireless paradigms that involve randomly located transmitters and receivers, such as femtocells and cognitive radios.

The goals of this article are:

- To argue that such tools are essential to understanding fundamental network qualities such as connectivity, coverage, and capacity
- To provide a basic primer on how these tools may be applied for modeling and analysis
- To identify future research areas where these tools may be useful, while enumerating what innovations in the theory itself would be helpful.

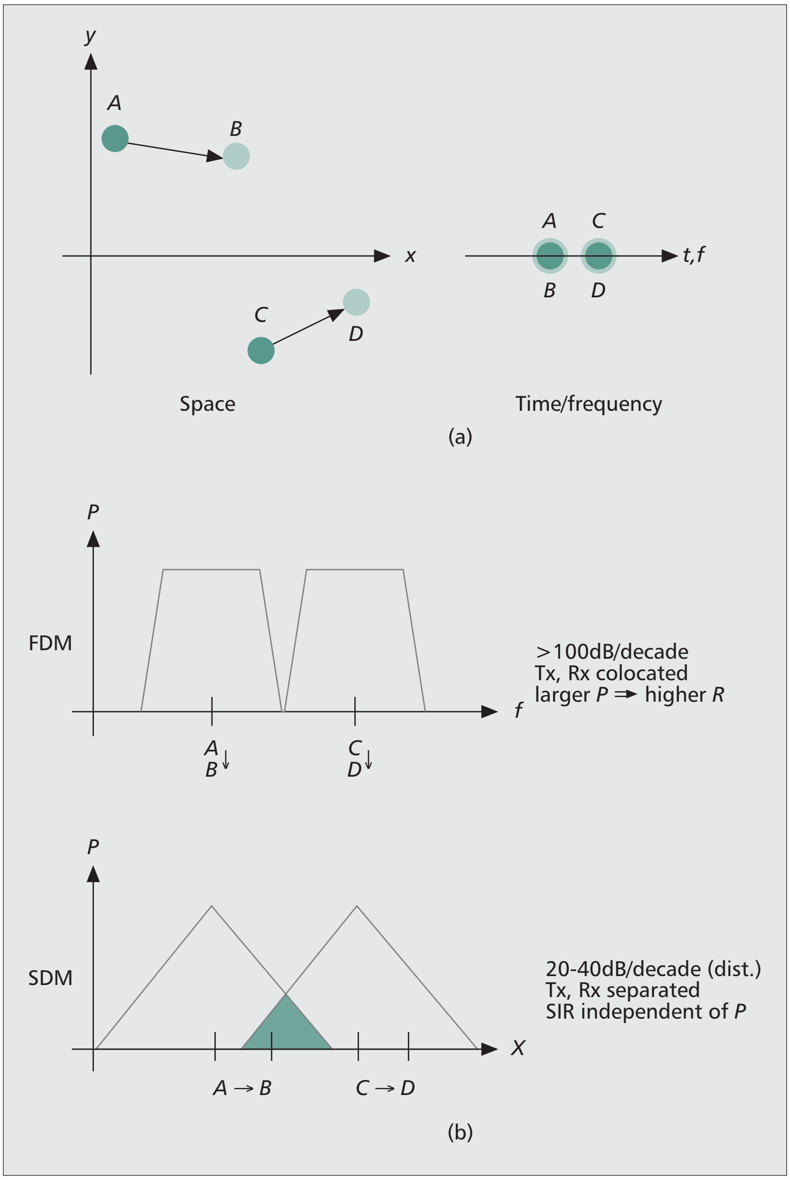

Figure 1. *a): Transmissions in time/frequency can be tightly packed, but not in space. (b) In frequency division multiplexing, power falls off very quickly outside the desired band. In space, the power decay is slow and uncontrollable, and causes potentially severe interference to neighbors.*

Spatial Models and Metrics for Wireless Networks

The effectiveness and efficiency of a wireless network are best characterized by metrics such as connectivity, capacity/throughput, and reliability (e.g., packet error rate or outage probability). Each of these metrics is a complicated function of all the *links* in the network that potentially connect each pair of nodes. As the network grows in size to even a moderate number (10–100) of nodes, the number of possible combinations of communicating pairs explodes and a deterministic evaluation of these metrics is cumbersome at best. In order to speak coherently about the performance of the network and to meaningfully compare various techniques and protocols, a sensible approach is to discuss statistical conditions over the network, such as averages or probability of outage/success.

Modeling Node Locations: Spatial Point Processes

Inter-node distances significantly affect network performance. Therefore, a stochastic model for the node locations (i.e., a *spatial point process*) is needed. Spatial point processes are the general-

Point Process	Key Properties	Practical Example	Reference
Poisson (PPP)	Mutual independence between (transmitting) node locations.	Ad hoc networks with pure random channel access.	Fig. 3. Most prior work.
Binomial	Similar to PPP as far as i.i.d. node locations, but with a fixed number of nodes in a given area.	A known number of relays or mobile users deployed at random in a cell of known size	[4]
Poisson cluster (PCP)	Clustering of nodes, with independence between cluster locations.	Sensor networks, military platoons, an urban network with dense hotspots.	[5]
Poisson plus Poisson Cluster	Independence between the PCP and the PPP. Attraction between nodes.	PPP represents the mobile users in a macrocell and the PCP represents femtocells or hotspots.	Fig. 5. [6]
Matern hard core	Minimum distance between nodes.	Carrier sensing wireless networks with collision avoidance, e.g. WiFi.	Fig. 3. [2]
Determinantal	Repulsion between nodes, e.g. Ginibre Process.	CSMA networks, networks with soft minimum distance.	Fig. 2

Table 1. *Common spatial models, in approximate order of simplicity/tractability.*

ization of point processes indexed by time to higher dimensions, such as 2-D and 3-D space. Stochastic geometry provides the tools to analyze important quantities such as interference distributions and link outages, and thus permits statistical statements about network performance [2, 3]. It also allows the designer to focus on a single receiver or link by making the notion of a *typical node* or a *typical link* mathematically precise. Due to its analytical tractability and practical appeal in situations where transmitters and/or receivers are located or move around randomly over a large area, the (homogeneous) Poisson point process (PPP) has been by far the most popular spatial model. For example, in the 2-D PPP, each node takes up an independent location characterized by a pair of coordinates (x_i, y_i), the density of nodes in a unit area is λ, and so the average number of nodes in an area A is λA. Finally, the probability that there are n nodes in A is given by the Poisson distribution and thus equal to $(\lambda A)^n e^{-\lambda A}/n!$

Recent work has also considered more general models such as cluster models — in which nodes tend to cluster in certain locations — or hard-core models, in which nodes have a guaranteed minimum separation, for example due to a carrier sensing MAC protocol that avoids nearby simultaneous transmissions [2]. More general point processes typically result in less tractable expressions that include integrals that must be numerically evaluated, which is still much simpler than an exhaustive network simulation. Some useful point processes for wireless network modeling are summarized in Table 1, and a few sample illustrations are given in Fig. 2.

SINR — The Building Block Metric

Each of the key metrics follows directly from the received signal-to-interference-plus-noise ratio (SINR) on one or more links, so understanding the SINR is essential. The SINR is the instantaneous ratio of desired energy to interference and noise energy, and so is a random variable that depends on many factors. The most important factors are as follows.

The Distance between the Desired Transmitter and the Desired Receiver — Based on electromagnetic laws, the desired received power falls off with distance and obeys an inverse power-law where the exponent is known as the *path loss exponent*. In free space the power decay is quadratic with distance, but over ground the path loss exponent is usually better modeled by a value between 2.5 and 4 because of scattering and absorption. The difference in received power between a 5 meter link and a 100 meter link is a factor of at least 400 (path loss exponent of 2), but is more likely 10,000 or more for more typical path loss exponents. This multiple order-of-magnitude attenuation and dynamic range is fundamental to the behavior of wireless networks.

The Set of Active Transmitters — There are many potential combinations of active transmitters in even a moderate sized wireless network. The set of active transmitters is often chosen by the MAC protocol. To each receiver, the other active transmitters appear to be interferers.[1]

The Sum Interference Power — The sum interference power depends on the set of interfering transmitters and their distances from each desired receiver. In networks of moderate to high density the interference power is usually much larger than the noise power.

The Noise Power — The impact of the ambient noise power on the SINR depends upon received signal and interference powers: under low transmission power the SINR is noise-limited, while under high transmission power the SINR is interference-limited.

Many other factors can affect the SINR including random propagation effects (fading and shadowing), specific transceiver design practices (for example, the use of multiple antennas or interference cancellation), and power control. But the spatial interactions are the most fundamental and inescapable — general fading, shadowing, and power control models can be (and

[1] *Using two or more separate frequency bands adds a degree of freedom for scheduling and would usually reduce the number of interferers per band, but in its essence the problem is unchanged in each band. Hence we restrict our attention in this article to operation in a single band.*

have been) added to the baseline model that is the focus of this article. A fairly general but simple mathematical description of the SINR at a typical node located at origin o is:

$$SINR_o = \frac{h_{oo}\rho_o r^{-\alpha}}{N_o + \sum_{i \in \Phi} \rho_i h_{io} |X_i|^{-\alpha}},$$

where h_{io} is the (power) fading coefficient of the channel to the desired receiver o from node i, ρ_i. is the transmit power of transmitter i, N_o is the noise power, and Φ is the set of interfering nodes (Φ is a subset of all possible transmitters). The desired transmitter is a distance r from the desired receiver, while the ith interferer is a distance X_i away. By drawing the distances according to a probabilistic spatial model, the randomness in locations along with many other basic aspects of the network (e.g., path loss) are consolidated into a single random variable, the SINR.

We now will briefly overview how the SINR may be used to specify and ultimately compute metrics of interest, namely the connectivity/coverage and the capacity/throughput. As shown below, it is possible (in fact preferable) to incorporate reliability into both of these classes of metric, so considering reliability separately is unnecessary.

Connectivity and Coverage

The connectivity of a random network can be described as the probability that an arbitrary pair of nodes are able to exchange information at a specified rate. For example, if this probability is 0.9 for a random selection of a source-destination pair, then one would say that the network is 90 percent connected. The minimum power requirement for wireless network connectivity is intimately connected with percolation thresholds; this formed the basis for many early results. In the simplest case of direct transmission, i.e., single hop communication, the probability of connectivity is simply Pr[SINR > β], where β is the minimum required SINR that is considered acceptable, and is a tunable parameter and SINR is the signal-to-interference and noise ratio of a *typical* link. Note that for a desired rate R in bits per second, $\beta \approx \Gamma(2^R - 1)$, where $\Gamma \geq 1$ is the SNR *gap* from Shannon rate signaling.

In many wireless networks of interest, a single hop is all that is required or in fact allowed (e.g., traditional cellular networks). In such cases, the region of connectivity around a given transmitter is known as its *coverage area*. More generally, a source and a destination may communicate using one or more intermediate relays, in which case a path through the network must be found where each hop has an SINR greater than β. Also, the case where there is only one active flow in the network (in which case there is no interference from nodes not participating in transmitting this flow) and the case where there are many flows, where each node may be serving as a relay for one or more flows, need to be distinguished. There are many ways to describe and quantify network connectivity, but at the core, they all require that individual pairs are able to communicate, which is dictated by the SINR.

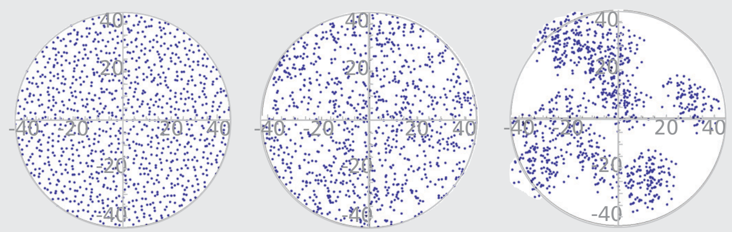

Figure 2. *Three sample point processes. Poisson-distributed nodes have independent locations, whereas a Ginibre determinantal process can be used to model more evenly distributed nodes, and cluster process can be used to model situations where nodes are likely to be close to one another, e.g. due to terrain.*

Throughput

Throughput is one of the most important performance metrics for wireless networks, and a number of different notions of throughput exist.

Link Throughput — Spatial models lend themselves to an analytical characterization of the per-link throughput, which is a critical determinant of end-to-end rate in multi-hop networks and is the quantity of interest for single-hop networks. Per-link throughput is dictated by the SINR, and can be defined in different ways. The average per-link throughput is $R_{avg} = E[\log(1 + SINR_{ij}/\Gamma)]$, where the average is with respect to the sources of randomness encapsulated in the random variable SINR (e.g., locations and fading). This metric can be appropriate for settings in which the transmitted rate is adjusted to the instantaneous SINR, whereas outage-based metrics are more appropriate when dynamic rate adjustment is not performed. The outage capacity of a link is the largest rate (or mutual information) that can be supported with a certain probability, for example 0.95, and so naturally includes reliability. In terms of SINR, this can be expressed in terms of a target outage probability ε as

$$C_{out} = \max \log(1 + \beta) : \Pr[SINR > \beta] > 1 - \varepsilon$$

For the network as a whole, it is then necessary to determine the outage Pr[SINR < β]. However, this depends on the Tx-Rx distance, and the locations of the active transmitters. Clearly, if fewer transmitters are active, then the SINR and hence the outage capacity can be increased, but the overall network throughput would also decrease. It is necessary to balance these two effects with a different metric. One such metric is the *transmission capacity*, first defined in [7] as

$$\tau_\varepsilon = (1 - \varepsilon)\lambda b$$

where λ is the maximum average number of active transmitters sending a rate of b b/s/Hz per unit area for which the outage probability is less than ε. In order words, the transmission capacity is the average number of successful active links of a certain rate that can be supported per square meter in the network [8]. It

For large wireless networks (e.g., ad hoc) where single hop communication is not possible, it is desirable to know the end-to-end rate that is supportable between a typical source-destination pair in the network. This is difficult to compute.

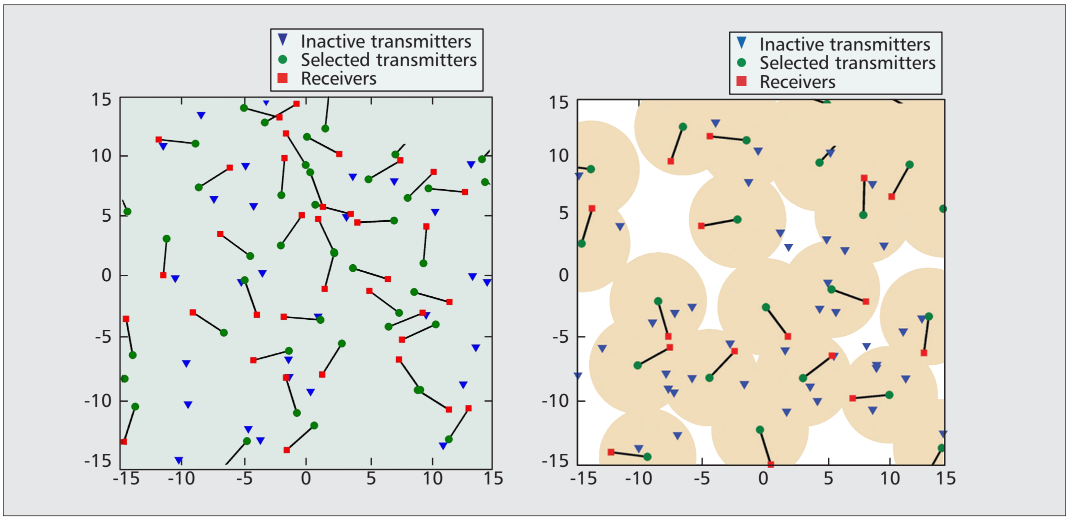

Figure 3. *(Left) A transmitter set resulting from an ALOHA MAC acting on node set initially distributed as a Poisson point process. (Right) A Matern hard-core process which models a CSMA/CA MAC. The discs represent an exclusion zone around each transmitter.*

has units of area spectral efficiency, for example b/s/Hz/m^2. As will be shown below, this metric is in fact computable over a wide range of network models.

End-to-End Rate — For large wireless networks (e.g., ad hoc) where single hop communication is not possible, it is desirable to know the end-to-end rate that is supportable between a typical source-destination pair in the network. This is much more difficult to compute because it depends on routing strategies, the retransmission strategies, and further depends on the reliability and rate of each hop. For certain strategies, however, end-to-end rate is a simple function of the per-link throughput and thus can be computed. Note also that end-to-end rate ties directly to the *transport capacity*, which is an end-to-end rate metric of units bit-meters/sec that incorporates distance and node locations, and gives credit in proportion to the distance the information is transported [9]. Measuring performance in terms of both achieved rate and distance traveled is also found in the effective forward progress metric used in early work in packet radio networks from the late 1970's. Because spatial models provide an explicit distribution on distances, such models are also amenable to quantification of transport capacity.

Applying Spatial Models

We now consider three types of wireless networks where spatial models play a central role: ad hoc networks, femtocells, and cognitive radio.

Ad Hoc Networks

Ad hoc networks — purely wireless networks in which all nodes in the network must exchange information with each other without any wired backhaul — are the framework in which spatial models have been most widely embraced. Indeed, the classical results on throughput scaling for ad hoc networks are fundamentally based on a spatial model and a spatial metric in which progress is measured in terms of rate times distance, as just discussed. Scaling laws do not reveal the effect of physical layer algorithms, channel access protocols, so now we consider how spatial models provide mechanisms to determine other fundamental properties of an ad hoc network.

The network in Fig. 3 (left), sometimes called a bipolar or dumbbell model, has each transmitter paired with a receiver in a random direction at a distance r. This model has received considerable attention, and the probability of the connectivity of an arbitrary link, assuming Rayleigh fading, relative to an SIR target β (noise neglected here for simplicity) is

$$P_s = \Pr\left[\mathrm{SIR} > \beta\right] = \exp\left(-\lambda r^2 \beta^{\frac{2}{\alpha}} C(\alpha)\right),$$

where λ is the density of transmitters, and $C(\alpha)$ is only a function of α [10]. It follows that the transmission capacity for an outage constraint ε is

$$\tau_\varepsilon = \frac{\varepsilon - 1}{r^2 \beta^{\frac{2}{\alpha}} C(\alpha)} \log(1-\varepsilon),$$

which gives the number of successful transmissions per unit area. If desired, τ_ε can be multiplied by $\log(1 + \beta/\Gamma)$ to give units of bits per symbol per square meter, or area spectral efficiency. The transmission capacity provides a clear view into how area spectral efficiency in a large ad hoc network depends on the basic network parameters. For example, the fact that transmission capacity decreases as r^2 provides a *sphere packing* interpretation in which each successful transmission consumes an area that depends upon the transmit-receive distance and the SINR threshold. This dependence upon the SINR threshold could, for example, then be used to find the threshold that maximizes area spectral efficiency.

The above result holds under the assumption that the set of active transmitters form a homogeneous PPP (the receivers are thus not a part of the underlying process). However, as noted earlier, more sophisticated spatial models can also be analyzed, such as assuming the active transmitters are distributed according to a hard-core process (emulating a CSMA/CA MAC) or that the transmitters and receivers are chosen from a common point process. More sophisticated transmission protocols can also be introduced fairly easily in the model by simply changing the starting SINR expression given in Eq. 1. For example, multi-antenna beamforming, spread spectrum, and power control can all be handled through appropriate modification of the SINR. An overview of these and other generalizations is given in [8]. This model has also been extended recently to a basic multihop model in [11].

Cognitive Radios and White Space

Scarcity of bandwidth and the allegedly sparse use of licensed spectrum by incumbents has led to the popularity of *cognitive* radios, which attempt to find locally unused spectrum and communicate over it opportunistically. The viability of this aggressive new approach to frequency reuse has been endorsed by the United States FCC in its 2009 Whitespace ruling, which lays down conditions under which cognitive radios can utilize previously licensed spectrum, namely the former analog TV bands, the majority of which are in the 470–806 MHz range. A primary consideration of a cognitive radio is the likelihood of interfering with a *primary*, or licensed, user of the spectrum. The probability of this occurring must be held small, and it clearly depends on the spatial density and the typically unknown locations of the primary receivers.

A cognitive radio network is depicted in Fig. 4. A simple model of a cognitive network consists of primary receivers and secondary transmitters/receivers distributed as homogeneous Poisson point processes Φ_p and Φ_s on the plane. Since the interference caused by the secondary transmitters at the primary receivers should be small, no secondary node close to an active primary user should be allowed to transmit. Hence the secondary transmitter set $\Phi_{st} \subset \Phi_s$ consists of only those secondary users whose distance is greater than some fixed value from every primary receiver. In practice of course, secondary nodes may not be able to ascertain precisely where the primary receivers are. In this case, they may attempt to ensure with high probability that the SINR at the secondary receiver is above a threshold, which is possible as long as they know just the primary receiver density. In a similar manner to the approach for ad hoc networks outlined above, outage probabilities and transmission capacities can be computed for each category of user [12, 13], and different communication protocols and algorithms can be evaluated to see how they affect the transmission capacity in each tier. Amongst other insights, this confirms that cognitive radios will typically cause outage to incumbent devices with positive probability; but this can be minimized by appropriate application of multiple antennas, sensing protocols, and other interference management techniques. There is a considerable amount ongoing research on various spatial models of different types of cognitive radio networks.

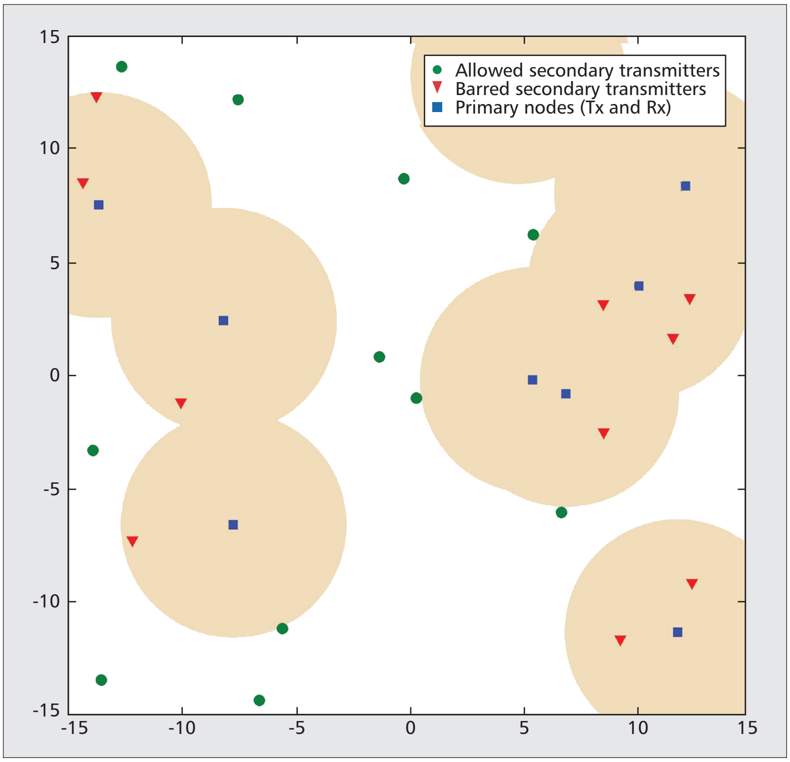

Figure 4. *The figure represents a simple point process model of a cognitive network where the primary nodes (transmitters and receivers) and the secondary transmitters are modeled as homogenous Poisson point processes of densities 0.01 and 0.025 respectively. Only secondary transmitters that are at least a distance 5 from any primary node are allowed to transmit. The secondary receivers are not shown.*

Femtocell Networks

Femtocells are very small, inexpensive base stations that overlay an existing cellular network to improve the capacity and coverage, particularly to indoor users. The femtocells can either be installed by end-users and companies, e.g., to improve their in-home and in-office coverage, or directly by the network operator, e.g., to improve capacity in airports, stadiums, and other areas of dense demand [14]. The market for femtocells is in an early phase but projected to reach 40 million units a year by 2013 according to an April 2010 New York Times technology article [15]. Naturally, tens of millions of arbitrarily-located devices that interfere with the carefully planned and deployed macrocell network is a source of serious concern for network operators. How can the potential large benefits of femtocell deployments be balanced with their potentially deleterious effect on the existing cellular network?

Stochastic geometry provides an essential toolkit for understanding femtocell deployments. A natural model for two-tier networks — consisting of tier 1 base stations and tier 2 femtocells — is to model the femtocell locations as a point process of density lf overlaying a regular grid of base stations. Shown in Fig. 5, the mobile users are also randomly located and can be mod-

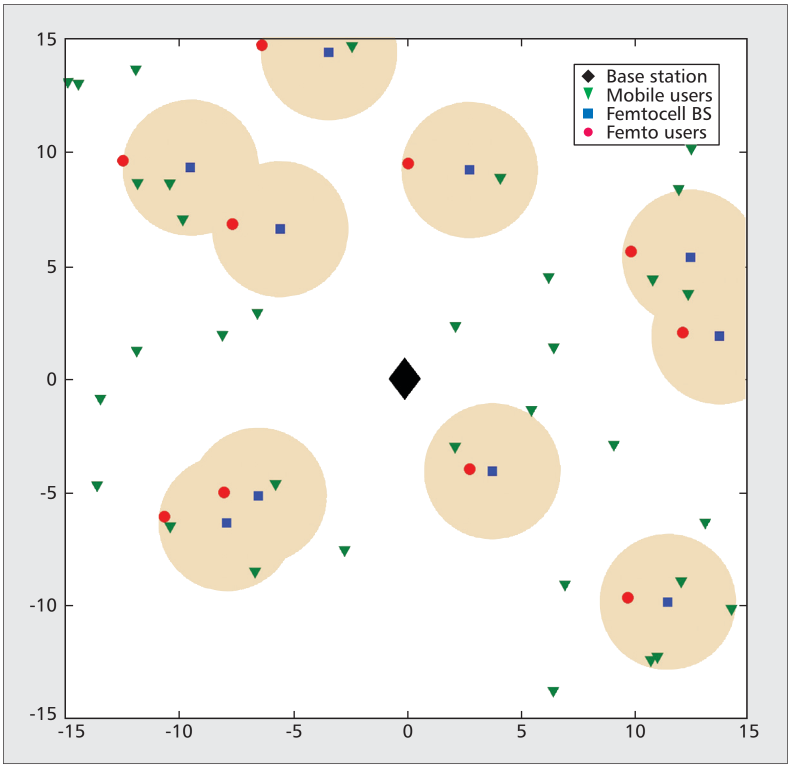

Figure 5. *A square cell overlaid by femtocells. The femto BS's are modeled by a Poisson point process (blue squares) and they serve a disc of radius 2. The green dots represent mobile users (non-femto cell users) which communicate with the main base station, which is represented as a black diamond. Note that it is possible for mobile users to be inside the range of a femtocell but not use the femtocell.*

eled as a point process of density l_c, where typically $l_c \gg l_f$. The interference at a given mobile user, for example, now consists of interference from neighboring base stations as well as from randomly placed femtocell base stations. Similarly, the interference at a femtocell base station is the aggregation of interference from all the uplink mobile users, and is typically dominated by a small number of mobile users transmitting at relatively high power up to the main base station. Again, each receiver's SINR can be carefully modeled using random spatial models for the interference from femtocells, mobile users, base stations, and femtocell users, as appropriate. The allowable density of mobile users can then be traded off with the femtocell density using outage probability or transmission capacity, and different techniques for cross-tier interference suppression and avoidance can be evaluated [6]. This is conceptually similar to a capacity region, where now the two competing axes are femtocell vs. macrocell achievable rate.

Future Directions for Spatial Models

Random spatial models will become increasingly relevant for the dense and complex wireless networks that will emerge over the next decade. This article has attempted to give a high-level view of the importance of such models and how they can be applied to different types of wireless networks. Much more work is needed, on the fundamental mathematics, confirmation of the accuracy of the models in the various scenarios, and further application of spatial models to emerging candidate protocols and networks.

The most tractable results from stochastic geometry tend to rely on a few assumptions that may not accurately hold in practice. The most important are the homogeneous Poisson point process for transmitting node locations and the neglect of temporal and spatial correlations. It is desirable to relax these assumptions in the coming years. There will be an ongoing trade-off between tractability and generality, but as random spatial models gain acceptance, it may be possible to improve the perceived tractability by identifying canonical solutions that may not be closed-form. For example, most probability of error expressions include an integral over Gaussian tail, accepted universally as the Q-function. The Q-function is not closed-form, but is often treated as such because it is so common. Similarly for complex spatial models, it may be necessary to simply define and name often-recurring integrals and live with them.

Even when analysis with random spatial models is not fully tractable, network performance can often be easily simulated because network-wide performance can be characterized by the performance of a typical node. Thus, agreed upon random spatial models will allow for standardized rapid benchmarking of wireless network protocols, just as well-accepted AWGN and fading channel models have been indispensable for fairly comparing techniques for point-to-point links.

Acknowledgments

The authors gratefully acknowledge feedback and suggestions from Angel Lozano, Robert Heath, Illsoo Sohn, Kostas Stamatiou, Jun Zhang, and the anonymous reviewers.

References

[1] D. C. Cox and H. Lee, "Physical Relationships: Exploring Fundamental Relationships Between Transmission Rate and Range for Wireless Systems," *IEEE Microwave Mag.*, Aug. 2008, pp. 89–94.
[2] F. Baccelli and B. Blaszczyszyn, *Stochastic Geometry and Wireless Networks, Foundations and Trends in Networking*, NOW Publishers.
[3] M. Haenggi and R. K. Ganti, "Interference in Large Wireless Networks," *Foundations & Trends Net.*, 2009.
[4] J. Zhang and J. G. Andrews, "Distributed Antenna Systems with Randomness," *IEEE Trans. Wireless Commun.*, Sept. 2008.
[5] R. K. Ganti and M. Haenggi, "Interference and Outage in Clustered Wireless Ad Hoc Networks," *IEEE Trans. Info. Theory*, Sept. 2009.
[6] V. Chandrasekhar and J. G. Andrews, "Uplink Capacity and Interference Avoidance for Two-Tier Femtocell Networks," *IEEE Trans. Wireless Commun.*, July 2009.
[7] S. Weber *et al.*, "Transmission Capacity of Wireless Ad Hoc Networks with Outage Constraints," *IEEE Trans. Info. Theory*, Dec. 2005.
[8] S. Weber, J. G. Andrews, and N. Jindal, "An Overview Of The Transmission Capacity Of Wireless Networks," *IEEE Trans. Commun.*, to be published.
[9] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Info. Theory*, Mar. 2000.
[10] F. Baccelli, B. Blaszczyszyn, and P. Mühlethaler, "An ALOHA Protocol for Multihop Mobile Wireless Networks," *IEEE Trans. Info. Theory*, Feb. 2006.
[11] J. G. Andrews *et al.*, "Random Access Transport Capacity," *IEEE Trans. Wireless Commun.*, June 2010.

[12] C. Yin *et al.*, "Transmission Capacities for Overlaid Wireless Ad Hoc Networks with Outage Constraints," *IEEE ICC*, Dresden, Germany, June 2009.

[13] K. Huang, V. Lau, and Y. Chen, "Spectrum Sharing between Cellular and Mobile Ad Hoc Networks: Transmission Capacity Trade-Off," *IEEE JSAC*, Sept. 2009, pp. 1256–67.

[14] V. Chandrasekhar, J. G. Andrews, and A. Gatherer, "Femtocell Networks: a Survey," *IEEE Commun. Mag.*, Sept. 2008.

[15] M. Richtel, "Bringing You a Signal You're Already Paying For," *The New York Times*, Apr. 6, 2010.

ADDITIONAL READING

[1] M. Haenggi *et al.*, "Stochastic Geometry and Random Graphs for the Analysis and Design of Wireless Networks," *IEEE JSAC*, Sept. 2009.

[2] D. Stoyan, W. Kendall, and J. Mecke, *Stochastic Geometry and its Applications, 2nd Ed.*, Wiley, 1996.

[3] S. Weber, J. G. Andrews, and N. Jindal, "The Effect of Fading, Channel Inversion, and Threshold Scheduling on Ad Hoc Networks," *IEEE Trans. Info. Theory*, Nov. 2007.

[4] J. F. C. Kingman, *Poisson Processes*, Oxford University Press, 1993.

[5] M. Z. Win, P. C. Pinto, and L. A. Shepp, "A Mathematical Theory of Network Interference and Its Applications," *Proc. IEEE*, Feb. 2009.

BIOGRAPHIES

JEFFREY G. ANDREWS [SM'06] (jandrews@ece.utexas.edu) is an Associate Professor in the ECE Department at UT Austin, where he is the Director of the Wireless Networking and Communications Group. He received a B.S. with High Distinction from Harvey Mudd College and his M.S. and Ph.D. from Stanford University. He has previously worked at Qualcomm and consulted for the WiMAX Forum, Microsoft, ADC, Apple, Clearwire, and NASA. He is co-author of Fundamentals of WiMAX (Prentice-Hall, 2007) and Fundamentals of LTE (Prentice-Hall, 2010). He received the NSF CAREER award in 2007 and is the Principal Investigator of a nine university team in DARPA's Information Theory for Mobile Ad Hoc Networks program. He has received best paper awards at IEEE GLOBECOM (2006 and 2009) and Asilomar (2008), the latter two for applications of stochastic geometry to femtocells, and the 2010 IEEE Communications Society Best Tutorial Paper award for "Stochastic Geometry and Random Graphs for the Analysis and Design of Wireless Networks."

RADHA KRISHNA GANTI [M'10] is a Postdoctoral researcher in the Wireless Networking and Communications Group at UT Austin. He received his B.Tech. and M.Tech. in EE from Indian Institute of Technology, Madras, and Masters in Applied Math and Ph.D. in EE from University of Notre Dame in 2009. His doctoral work focused on the spatial analysis of interference networks using tools from stochastic geometry. He is co-author of the monograph Interference in Large Wireless Networks.

MARTIN HAENGGI [SM'04] is an Associate Professor in the EE Department at the University of Notre Dame. He received the M.S. and Ph.D. degrees from the Swiss Federal Institute of Technology in Zurich, Switzerland (ETHZ) in 1995 and 1999, respectively. He received the ETH Medal for both his M.Sc. and Ph.D. theses and a CAREER award from the U.S. National Science Foundation in 2005. He served as the Lead Guest Editor for an issue of the IEEE Journal on Selected Areas in Communications on stochastic geometry and random graphs for wireless networks in 2009 and currently is an Associate Editor of the IEEE Transactions on Mobile Computing and the ACM Transactions on Sensor Networks. He co-authored the monograph Interference in Large Wireless Networks (NOW Publishers, 2009), and received the 2010 IEEE Communications Society Best Tutorial Paper award for "Stochastic Geometry and Random Graphs for the Analysis and Design of Wireless Networks".

NIHAR JINDAL [M'04] is an assistant professor in the ECE Department at the University of Minnesota. He received his B.S. from U.C. Berkeley, and his M.S. and Ph.D. from Stanford University. He is currently an Associate Editor for the IEEE Transactions on Communications, and has also served as a guest editor for the EURASIP Journal on Wireless Communications Networking. He received the NSF CAREER award in 2007 and the University of Minnesota Guillermo E. Borja Award in 2010. He also received the IEEE Communications Society and Information Theory Society Joint Paper Award in 2005, and the Best Paper Award for the IEEE Journal on Selected Areas in Communications (IEEE Leonard G. Abraham Prize) in 2009.

STEVEN WEBER [M'03] is an Associate Professor in the ECE Department at Drexel University. He received the B.S. from Marquette University in 1996, and the M.S. and Ph.D. from The University of Texas at Austin in 1999 and 2003 respectively. His research interests are centered around mathematical modeling of computer and communication networks, specifically streaming multimedia and ad hoc networks.