

Ad Hoc Networks: To Spread or Not to Spread?

Jeffrey G. Andrews, University of Texas at Austin

Steven Weber, Drexel University

Martin Haenggi, University of Notre Dame

ABSTRACT

Spread spectrum communication — often called code-division multiple access — has been widely adopted over the years for many types of interference-challenged wireless communication systems including cellular and cordless telephones, wireless LANs and PANs, military applications, and global positioning systems. In this article we explore whether CDMA, in either its frequency hopping (FH) or direct sequence (DS) form, is an appropriate design approach for wireless ad hoc, or mesh, networks. One goal of this article is to help provoke a debate by explaining the main advantages and disadvantages of CDMA in the context of ad hoc networks as exposed by recent research. We argue that CDMA does not inherently improve the spectral efficiency of ad hoc networks; on the contrary, its valued interference averaging effect is not appreciable in ad hoc networks due to the irregular distribution of both the transmitters and receivers. On the positive side, both types (FH and DS) of spread spectrum allow for longer hop distances and a reversal of the usual relationship where the desired transmitter must be closer to the receiver than interfering transmitters. These two facts allow for significant advantages over narrowband systems in terms of energy efficiency and end-to-end delay.

INTRODUCTION

Applications of wireless ad hoc networks have expanded in recent years to include not only numerous military applications and emerging wireless sensor networks, but also many other exciting and commercially viable applications including wireless community broadband access, backhaul for wireless LAN access points, and range extension for cell-based networks [1]. Despite this high level of interest and commercial potential, many basic ad hoc network design principles are still not well understood, and one important design question is the focus of this article: Does it make sense to use spread spectrum in ad hoc networks? In this article an ad hoc network implies communication without the assistance of wired infrastructure. Naturally,

such networks will often have at least some connections to wired infrastructure, but we neglect this for simplicity of discussion.

Spread spectrum transmission has long been considered attractive for ad hoc networking for a number of reasons, including security and interference robustness [2, 3]. In this article we carefully scrutinize the supposed advantages of using spread spectrum — also known as code-division multiple access (CDMA) — in ad hoc networks. Similar to the highly contentious CDMA vs. time-division multiple access (TDMA) debate for cellular systems, the considerations for ad hoc networks are also laden with subtleties. In cellular networks, despite CDMA's apparent inferiority due to intentional self-interference, the exploitation of voice activity, frequency reuse, and fast power control were central to the ultimate success of CDMA. Analogously for ad hoc networks, it is crucial to adopt a network-level point of view that includes considerations such as network capacity, end-to-end delay, energy efficiency, channel access, and routing. However, the key traits of CDMA in an ad hoc network are very different than in cellular networks with centralized transmitters (downlink) and receivers (uplink). We now summarize the key traits of CDMA in ad hoc networks in terms of pros and cons, which are justified in detail in the body of the article.

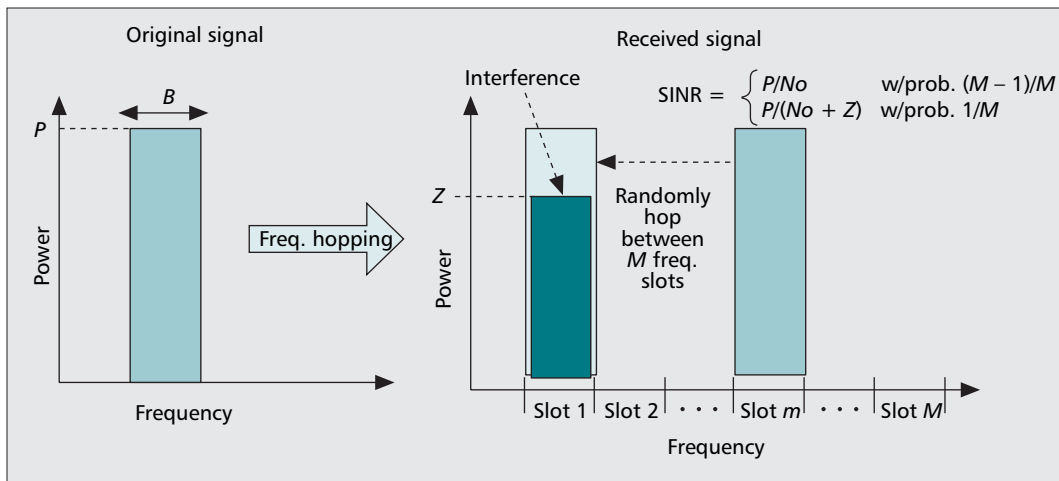
THE ADVANTAGES OF CDMA IN AD HOC NETWORKS

The advantages of CDMA in ad hoc networks are quite different than in cellular networks, and can also be distinct for the two different types of CDMA, frequency hopping (FH) and direct sequence (DS). FH and DS are described in more detail later.

Longer hops. DS and FH both allow for longer hops to be undertaken for a given network density. This allows more direct routing, reduced end-to-end delay, and perhaps counter-intuitively, reduced energy consumption.

Capacity enhancements. Neither FH nor DS increases the overall ad hoc network capacity on its own, in fact the opposite is true for DS. However, DS allows for the possibility of capacity-increasing interference canceling receivers, which are ineffec-

This research was supported by NSF grant no. 0635003 (Weber), no. 0634979 (Andrews), and the DARPA IT-MANET program, grant no. W911NF-07-1-0028 (all authors).



■ **Figure 1.** Frequency hopping works by randomly picking one of M frequency slots. A narrowband signal of similar bandwidth is usually avoided as M increases.

Spread spectrum uses noise-like code sequences to effectively increase the bandwidth to be far greater than the signal bandwidth. When spread spectrum is used to support multiple users, it is called CDMA.

tive in ad hoc networks unless DS is used. Both FH and DS also provide considerable frequency diversity, which helps overcome narrowband fading.

Network efficiency. The ability of DS-CDMA systems to successfully operate under low signal-to-interference-plus-noise ratio (SINR) permits communication with a larger number of potential interferers, which simplifies network coordination.

Security. Spread spectrum radios have innate security features: they are harder to jam, they make eavesdropping more difficult, and their presence is more difficult to detect. Although important for some applications, these topics are beyond the scope of this article.

THE DISADVANTAGES OF CDMA IN AD HOC NETWORKS

A common drawback of CDMA in cellular and ad hoc networks is that the system bandwidth needs to be considerably larger than the (per user) symbol rates. In ad hoc networks CDMA has two other important drawbacks:

Interference averaging is ineffective. Interference averaging, the hallmark of both DS- and FH-CDMA in cellular networks, does not pay off in ad hoc networks. The key reason is the lack of a centralized receiver and the associated power control to that receiver. Global power control is impossible in ad hoc networks; instead, usually just one or perhaps two interfering nodes dominate the interference power, which makes interference *avoidance* (via scheduling or slow FH) far more effective than interference *averaging* (using DS or fast FH to proportionally reduce the interference level).

Considerable setup costs. CDMA requires that both the transmitter and receiver have knowledge of:

- Agreed upon spreading (DS) or hopping (FH) sequences
- The current position in the sequence (i.e., time synchronization)

Acquisition of these is a nontrivial resource-consuming process, and unless the cost of code acquisition and synchronization is amortized over time, the above “pros” of CDMA may not justify this overhead. In practice this raises questions about

CDMA’s viability in ad hoc networks with mobility or bursty traffic, since these scenarios require frequent code acquisition and synchronization.

CDMA: A MODERN OVERVIEW


INTERFERENCE-LIMITED NETWORKS AND THE RECENT SUCCESS OF CDMA

Dense wireless networks are by nature interference-limited, which means that increasing the transmit power of all nodes in the network simultaneously will not substantially increase the overall throughput of the network. Ad hoc networks pose a particularly challenging interference environment because the lack of agreed upon centralized transceivers means that each receiver in the network must bound the level of interference in its vicinity to successfully receive the desired transmission.

Spread spectrum uses noise-like code sequences to effectively increase the bandwidth to be far greater than the signal bandwidth. When spread spectrum is used to support multiple users, it is called CDMA. The central tenet of CDMA is that designing for time or frequency orthogonality (as in TDMA or FDMA) is not appropriate, since neighboring (i.e., intercell) interference and other imperfections would compromise the orthogonality anyway. On the other hand, CDMA tolerates all sources of interference within bounds determined by the spreading gain. Due to its robustness, system capacity, and other implementation and political/economic factors, CDMA overcame extreme levels of early skepticism to become the underlying physical layer technology and multiple access scheme for all three important third-generation cellular standards: cdma2000, wideband CDMA (WCDMA), and TD-SCDMA. Based on this success, it is natural to seriously consider the viability of CDMA for the emerging class of ad hoc and mesh networks.

FREQUENCY HOPPING AND DIRECT SEQUENCE CDMA

CDMA techniques have historically been divided into two very different types of modulation: frequency hopping and direct sequence. In this arti-



■ **Figure 2.** Direct sequence works by spreading the signal over a larger bandwidth. After processing, the desired narrowband signal re-emerges while other interference is attenuated by a factor of M .

cle CDMA without further qualification refers collectively to both of these techniques. FH is depicted in Fig. 1. The total bandwidth W is divided into M frequency bands of bandwidth $B = W/M$. At each hop, the transmitter chooses one of the M bands based on a pseudorandom code sequence that is also known to the receiver. Assuming the transmitter and receiver are synchronized, they both hop in unison and are able to successfully communicate. If there are other users in the network, there are occasional collisions when two transmitters pick the same frequency band, but by coding over time, it is possible to recover from a moderate number of collisions. Fast hopping refers to hopping on the order of a symbol time, whereas slow hopping refers to hopping on the order of a packet time. Examples of well-known systems that use FH include Bluetooth, which has 80 frequency bands of 1 MHz width ($M = 80$, $W = 80$ MHz), and a hop interval of 625 μ s, and GSM (which is also TDMA), which has a variable number of possible frequency bands of width $B = 200$ kHz and a hop interval of 4.617 ms.

Direct sequence, shown in Fig. 2, also involves synchronized pseudorandom codes, but in this case a code sequence with bandwidth $W = B \cdot M$ is multiplied with the user's data sequence of bandwidth B , creating a transmitted sequence of bandwidth W . M is called the spreading factor. By correlating the same code with the received signal, the desired signal is converted back to a narrowband signal (i.e., bandwidth B), while the noise and interference stay at bandwidth W and hence are attenuated by a factor of approximately M at detection.

In summary, FH systems avoid interference with increased probability as the number of frequency slots M grows, whereas DS systems suppress interference by a factor of M . Both FH and DS entail some important design considerations relative to narrowband transmissions. First, the transmitter and receiver need to be synchronized and aware of each other's code sequences. Bluetooth provides a useful practical example of how an FH-CDMA ad hoc network can achieve both time and code synchronization using a straightforward paging and inquiry procedure. In a mobile ad hoc network where this procedure must be done frequently, this overhead may nonetheless be considerable. A second distinguishing factor of DS-CDMA in cellular systems

is its reliance on accurate power control. In ad hoc networks, however, equal received powers are impossible due to the random node positions and the distributed nature of the networks.

Despite their significant differences, DS- and FH-CDMA have many similar properties in power-controlled cellular systems and achieve comparable SINRs for the same system load. Both effectively "average" interference so that a system can be designed for the average, rather than worst case, interference. DS is generally preferred for multiuser cellular systems since it has smoother interference averaging properties (the received SINR does not fluctuate as much), easily allows for coherent modulation that gives a 3 dB gain, and, perhaps most important, allows strong error correction coding without sacrificing spreading gain. *In contrast to cellular systems, though, FH and DS have very different characteristics when used in an ad hoc network, and the trade-offs between FH, DS, and narrowband signaling are quite different.*

THE KEY FEATURES OF CDMA AD HOC NETWORKS

The fundamental metrics of interest in an ad hoc network are capacity, end-to-end delay, and energy efficiency; these often compete with each other. In this and subsequent sections we discuss how the distinctions between spread spectrum and narrowband ad hoc networks affect these performance metrics.

CAPACITY OR THROUGHPUT

"Capacity" is a suspect metric in an ad hoc network, since it is interdependent with delay, transmit distance, mobility, scheduling, and higher-layer network functionality. To ground the discussion, we consider the *transmission capacity*, which is the average number of reliable simultaneous links that can be active in a unit area under a specified typical outage probability. Transmission capacity is a single-hop capacity metric that permits precise throughput characterization, whereas the more commonly used metric of transport capacity [4] is a multihop metric that permits only an asymptotic scaling law of network throughput. It is reasonable to expect that spread spectrum would allow a higher transmission capacity C to be tightly upper bounded as [5]

$$C^{FH} < \frac{\epsilon}{\pi r^2} \frac{M}{\beta^{\frac{2}{\alpha}}}, \quad C^{DS} < \frac{\epsilon}{\pi r^2} \left(\frac{M}{\beta}\right)^{\frac{2}{\alpha}} \quad (1)$$

for a transmit distance of r , outage probability ϵ , path loss exponent $\alpha > 2$, spreading factor M , and target SINR β . Here we have assumed that the hopping speed is equal to the packet length (i.e., the duration over which outage is determined), the frequency-time slots are orthogonal, signal power decays with distance as $d^{-\alpha}$, and all active nodes are separated by at least a few wavelengths (this prevents anomalous behavior of the $d^{-\alpha}$ expression for small d).

Insights. Transmission capacity allows for a remarkably simple expression, which indicates that FH is better than DS by a factor of $M^{1-2/\alpha}$ (e.g., by a factor of \sqrt{M} for the typical value $\alpha = 4$). Similarly, one could ask if the bandwidth spent on spread spectrum is justified by the interference savings. For DS, the answer appears to be no: the capacity grows sublinearly with M (as $M^{2/\alpha}$) while the consumed bandwidth grows as M . For FH, there appears to be no fundamental bandwidth penalty since both the capacity and bandwidth increase linearly with M . The basic conclusion is that interference avoidance (by hopping) is preferable to interference suppression (by despreading) in an ad hoc network. This is a byproduct of the near-far problem, which is why the gain from interference avoidance increases for large α .

Caveats. There are a few caveats that should be noted before concluding that DS-CDMA has inferior capacity in an ad hoc network. First, the above results implicitly assumed ALOHA-type medium access control (MAC), that is, the transmitter locations are random and independent of one another. A better MAC for DS-CDMA, as seen below, deliberately clusters transmitters and receivers. Second, a matched-filter receiver was assumed, which achieves a SINR gain of M , which is known to be highly suboptimal (in theory) relative to a multiuser interference-canceling receiver. Finally, bandwidth efficiency may not be the key concern in all applications. In ultra wideband (UWB) communication, robustness to interference may be more important than absolute bandwidth efficiency. Nevertheless, it should be conceded that this initial evidence suggests DS-CDMA is not as promising for ad hoc networks as it was for cellular.

THE INTERFERENCE RANGE

The *interference range* is a common concept in understanding ad hoc networks; it is defined as the minimum distance s such that any interfering node within s of a receiver will by itself generate sufficient interference power to cause an outage at that receiver. Conceptually, it is simplest to think of the interference range as a disk around the desired receiver, as shown in Fig. 3, but in reality it is an irregular contour due to random channel effects. We call interferers within the interference range *dominant*; it can be shown that, to a first order, many performance metrics of interest may be obtained through consideration of dominant interferers alone. The interference radii under FH and DS are

$$s^{DS} = r \left(\frac{\beta}{M}\right)^{\frac{1}{\alpha}}, \quad s^{FH} = r\beta^{\frac{1}{\alpha}}, \quad (2)$$

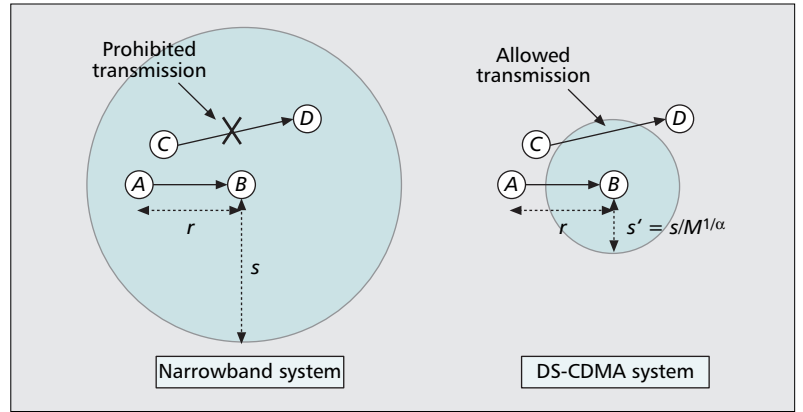


Figure 3. The interference range (s) and region (shaded) for narrowband and DS-CDMA. For NB, if $A \rightarrow B$ is scheduled, $C \rightarrow D$ is not allowed since it will cause an outage at B . In CDMA with sufficient spreading factor M , $C \rightarrow D$ is allowed since the interference range shrinks as $1/M^{1/\alpha}$.

where β is the required signal-to-interference ratio (SIR) and r is the distance separating the receiver from its associated transmitter. Note that the effective SIR requirement under DS is β/M , while the interference radius is independent of M for FH. However, in FH dominant interferers are removed with probability $(M-1)/M$.

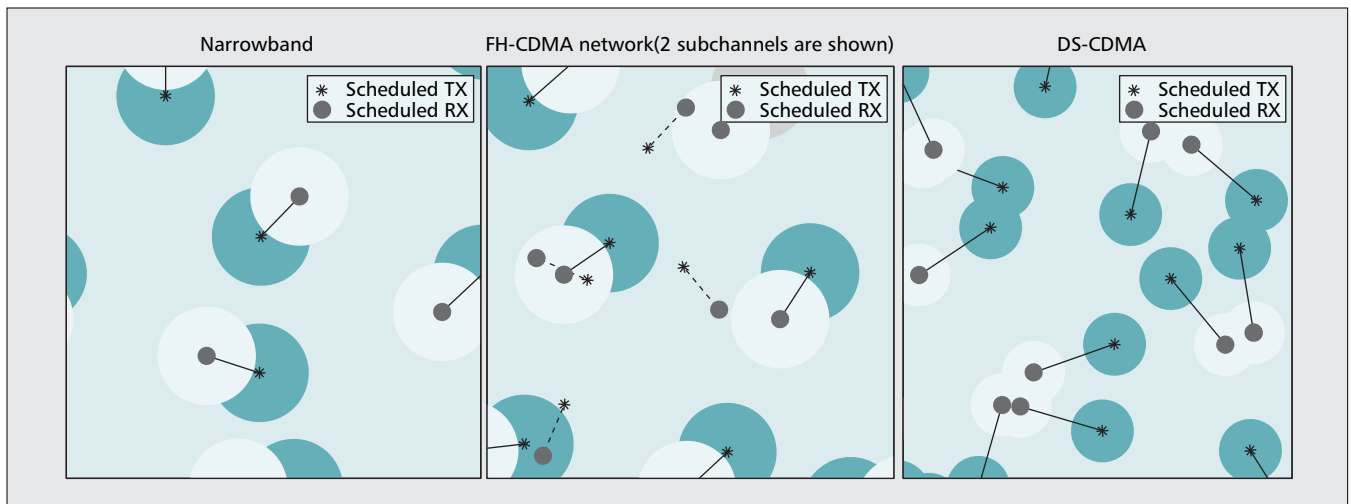
When nodes are uniformly spread out in space, it is possible to compute the probability of outage due to the presence of one or more dominant nodes. Setting this outage probability for both DS and FH equal to ϵ and solving for the radius r we obtain:

$$r^{DS} = cM^{\frac{1}{\alpha}}, \quad r^{FH} = c\sqrt{M}, \quad c = \beta^{\frac{1}{\alpha}} \sqrt{\frac{-\log(1-\epsilon)}{\pi\lambda}}, \quad (3)$$

where λ is proportional to the number of nodes in the network. This expression can be interpreted as saying that the maximum distance separating successful transmitter-receiver pairs scales in the spreading factor like $M^{1/\alpha}$ under DS and like \sqrt{M} under FH. Given $\alpha > 2$, this result demonstrates that FH allows longer hops to be taken than DS or narrowband, all else being equal. That said, the expression for s^{DS} demonstrates that $s^{DS} < r$ for sufficiently large M , that is, successful transmission under DS requires coordination only with nodes that are closer to the receiver than the transmitter. A key difference between spread spectrum and narrowband transmission can be observed from the above relations: with sufficient spreading gain, **the transmission range can be greater than the interference range**. In contrast, the interference range is larger than the communication range in both narrowband and FH, assuming $\beta > 1$, which is invariably true for all but the lowest conceivable data rates. In short, spread spectrum transmission has the interesting property of changing the usual relationship of $s > r$ to $s < r$, which has a number of important implications to be highlighted shortly.

MEDIUM ACCESS CONTROL FOR SPREAD SPECTRUM NETWORKS

Given these unique traits of spread spectrum — its superior allowable transmission density and increased interference-to-communication-range



■ **Figure 4.** Sample transmit/receiver pairs with near-optimal scheduling in narrowband, FH and DS-CDMA networks, from [7]. Narrowband systems require isolated Tx-Rx links. FH (center) allows some colocated links: disks are shown around the users in subband 1, no disks are around users in subband 2. DS throughput is maximized when transmitters and receivers cluster together, as seen at right.

ratio — it is clear that spread spectrum MAC should be designed differently than for a narrowband (NB) system. On the negative side, spread spectrum systems are burdened with considerable overhead in terms of exchanging and synchronizing their code sequences.

Frequency Hopping MAC Design — By its very nature, FH provides a high probability of interference avoidance, which is why it achieves high capacity even in dense networks. Thus, once the transmitter and receiver have acquired each other and are synchronized, the channel access part of the FH MAC is extremely simple since it need not perform contention resolution among colocated concurrent transmitters. One way to further improve the spectral efficiency of FH is to use adaptive spectrum sensing (similar to cognitive radio) or adaptive FH, where the hopping sequences of nearby nodes are acquired to preemptively avoid collisions. The trade-off incurred for these benefits is that before any communication can take place, the channel hopping sequence and present state must be agreed on, which consumes significant time and bandwidth resources.

Direct Sequence MAC Design — Because nearby interferers cause very strong interference, simply attenuating this interference by M is not particularly effective, as observed in Eq. 1. In cellular networks this problem can be mitigated using centralized power control, but in ad hoc networks power control is impractical since:

- There is no centralized authority to coordinate the required power levels.
- Even *perfect* power control would not preclude excessive interference at some receiving nodes due to the network geometry.

This necessitates contention resolution in both DS and NB ad hoc networks. This irreconcilable strong interference problem is a frequent criticism against DS-CDMA in ad hoc networks, but this is a misconception since NB systems suffer even more drastically from nearby interferers [6]. The key point is that both DS and NB systems require scheduling or contention resolution

since the use of quasi-orthogonal code sequences is insufficient to suppress nearby interferers.

The most popular contention resolving MAC is carrier sense multiple access (CSMA), but its popular implementation in wireless networks is highly spectrally inefficient, especially for DS, since it inhibits nodes around both the transmitter (which does the sensing) and the receiver, which responds to a request to send (RTS) message with a clear to send (CTS). An efficient MAC would only inhibit transmissions near the receiver. In fact, as shown in Fig. 4, clustered transmitters and receivers are highly desirable in a CDMA ad hoc network, and an optimal DS MAC will result in a large degree of clustering [7]. Therefore, a better approach for a DS MAC is for the receiver to instruct neighbors in its vicinity to suppress their transmissions. In contrast to NB systems, this explicit coordination is straightforward in DS systems, since the communication range is longer than the interference range. Therefore, *a DS receiver can communicate with all its potential strong interferers*. This large advantage again must be traded off with the overhead inherent in code acquisition and synchronization, which, as in FH, is particularly burdensome in a mobile network. Typically, DS code acquisition is achieved with progressively more wideband pilot signals, although another option is to use a separate NB control channel for this purpose. One considerable drawback of the NB control channel is that it is subject to NB fading, interference, and all the other impediments that motivate spread spectrum in the first place.

THE ADVANTAGES OF LONGER HOPS

It is commonly assumed that capacity is maximized and energy consumption minimized by routing through the nearest neighbors. The argument is that shorter-range transmissions cause less interference, permitting better spatial reuse. For example, if the transmission range for each hop in a route were reduced by half, the effective area of interference would decrease by $2^2 = 4$, while the number of hops would only double.

Hence, the overall network capacity would increase if all nodes reduced their transmit range. A further supposed advantage of nearest-neighbor routing is its improvement in energy efficiency. For a path loss of d^α , where d is the link distance, halving the transmission range would reduce the per-hop energy consumption by 2^α and the total energy over the two hops by $2^{\alpha-1}$.

While in some cases it may indeed be beneficial to route through the nearest possible neighbors, the above two arguments are simplistic in the face of important considerations regarding capacity, energy consumption, and end-to-end delay. As was argued in the previous section, DS spread spectrum allows longer hop ranges by a factor of $M^{1/\alpha}$, and FH by a factor of \sqrt{M} . Since an understanding of the effect of hop and route lengths is central to a debate on CDMA's merits, we now apply spread spectrum to some of the arguments of [8] and explain the key reasons that long hops are often preferable, which is a potentially significant advantage of spread spectrum over NB transmission.

Capacity — Although increased capacity is ostensibly the largest advantage of short hops, it is not at all clear that nearest-neighbor routing actually results in the best observed network throughput. Shorter hops at low power result in more transmissions, and hence lower interference levels, but for longer time periods. On the other hand, if all active transmitters increase their transmit power by a constant factor, the link SINRs can only increase, since the noise term becomes negligible. Hence, it is not obvious that many low-power transmissions are always superior to fewer high-power transmissions. In order to maintain a given end-to-end data rate (spectral efficiency), the per-hop rate needs to increase as the number of hops increases. This leads to an optimum number of hops that is considerably less than the maximum number of hops [9].

Energy Efficiency — We give three reasons that long hops are preferable from an energy perspective. First, the logic that shorter hops require less transmit power and hence reduce power consumption is suspect, since this assumes that reducing the distance by a factor x reduces power consumption at the transmitter by x^α . This is an oversimplified energy model; from a power amplifier efficiency standpoint it is far preferable to send at the maximum linear operating point and route as far as possible [10]. Backing off from this power level does not significantly reduce the current drawn by the power amplifier. Also, reducing the transmit power implies additional overhead due to power control and variable power transmission. Second, nearest neighbor routing requires many nodes to perform routing when they could otherwise go into sleep mode (which consumes 1 percent or less the power of being "awake"). Third, nearest neighbor routing reduces the network's ability of load balancing, since most nodes need to participate in the routing process.

Delay — The fact that short-hop routing incurs more delay than long-hop routing is not disputed. It is typically assumed that delay is proportional to the number of hops. Even this is optimistic, since employing many short hops:

- Increases the chance of packet loss
- Makes bottlenecks or *traffic jams* more likely to occur
- Increases the delay variance, making delay guarantees difficult

Longer hops also reduce the time for route discoveries; this advantage becomes more significant as node mobility increases.

In summary, since spreading increases the allowable hop distance for the same network density, FH and DS both appear to be very attractive means of increasing energy efficiency and reducing end-to-end delay. Even if nearest neighbor routing may turn out to maximize the throughput in many configurations, there may be many scenarios where delay and energy consumption are more pressing requirements. CDMA ad hoc networks thus provide an additional means of adaptation; for a modest sacrifice in spectral efficiency, FH and DS allow longer hops, so with the need for decreased power or delay (i.e., low batteries, real-time applications), spreading can be of pivotal assistance.

IMPROVING AD HOC NETWORK CAPACITY

A fundamental difference between DS-CDMA and both FH and NB is that the latter techniques are not designed to tolerate any co-channel transmissions, but instead avoid them by either scheduling (NB) or hopping (FH). DS-CDMA can also benefit substantially from scheduling, as well as from sophisticated receivers that cancel co-channel interference. In this section we explore the possible gain from each of these techniques, and see that in both cases DS systems are better suited to take advantage of their gains.

ADVANCED RECEIVERS

Multuser receivers have been widely studied in academia, and the principal approaches are well summarized in [11]. However, these receivers have never really caught on in industry, primarily due to their complexity for large numbers of users, incompatibility in actual wireless channels, and adversarial relationship with error correction codes [12]. Multuser detection is actually more attractive in ad hoc networks than in cellular systems due to the large benefit attainable from canceling just a few interfering nodes. Interestingly, *only DS-CDMA systems stand to benefit from most achievable multuser detection techniques*, since only interferers stronger than the desired signal can typically be cancelled. In DS systems, this is sufficient since all interferers (in particular those farther away than the desired transmitter) are attenuated by the spreading factor, and all those strong enough to cause an outage are within communication range, as discussed earlier. However, in NB or FH systems, these further interferers can still cause an outage, but since they are outside of communication range, it is nearly impossible to cancel them.

Successive interference cancellation (SIC) is an appropriate type of multuser detection for ad hoc networks, given its theoretical optimality and amenability to implementation when the number of nodes to be cancelled is small. The transmission capacity of imperfect SIC is shown in Fig. 5

In summary, since spreading increases the allowable hop distance for the same network density, FH and DS both appear to be very attractive means of increasing energy efficiency and reducing end-to-end delay.




Figure 5. SIC's effect on transmission capacity. As the number of cancellable nodes K increases, the additional gain from SIC is essentially zero, even with very accurate interference cancellation, that is, small ζ , where ζ is the fraction of interference power left for each user after cancellation.

vs. the number of cancelled nodes [13]. Some key conclusions, which apply to any type of imperfect interference suppression in ad hoc networks, are:

- Similar to cellular, perfect SIC increases the capacity by perhaps an order of magnitude. *Good news: large potential throughput gain.*
- Unlike cellular, most of the gain is achieved by just canceling the one or two dominant interferers. *Good news: low complexity and latency.*
- Unlike cellular, the interference cancellation is exceptionally sensitive to the amount of residual interference. *Bad news: channel estimation must be extraordinarily accurate.*

The key fact is that the residual interference of the strongest interferer is usually more important than the full interference of the other interferers. Residual interference is inevitable, and results primarily from imperfect estimates of the channel amplitude and phase, which is required for reconstructing the signal to cancel [12]. This is why even for relatively accurate interference cancellation (in practice 90–95 percent of the interference cancelled for each node would be considered a success, i.e., $\zeta = 0.05$ – 0.1), there is no discernible benefit to canceling more than 1 node. Therefore, although the large potential benefits of multiuser detection are unique to DS ad hoc networks, designers should approach claims of huge gains with skepticism.

SCHEDULING

MAC scheduling is critical to the efficient operation of DS-CDMA ad hoc networks. The key difference between DS scheduling and NB and FH scheduling is that due to DS's interference suppression margin, receivers should be clustered close to each other, as should transmitters, as discussed earlier. Although we previously discussed the interference range of DS ad hoc networks, we did not quantify the optimum *guard*

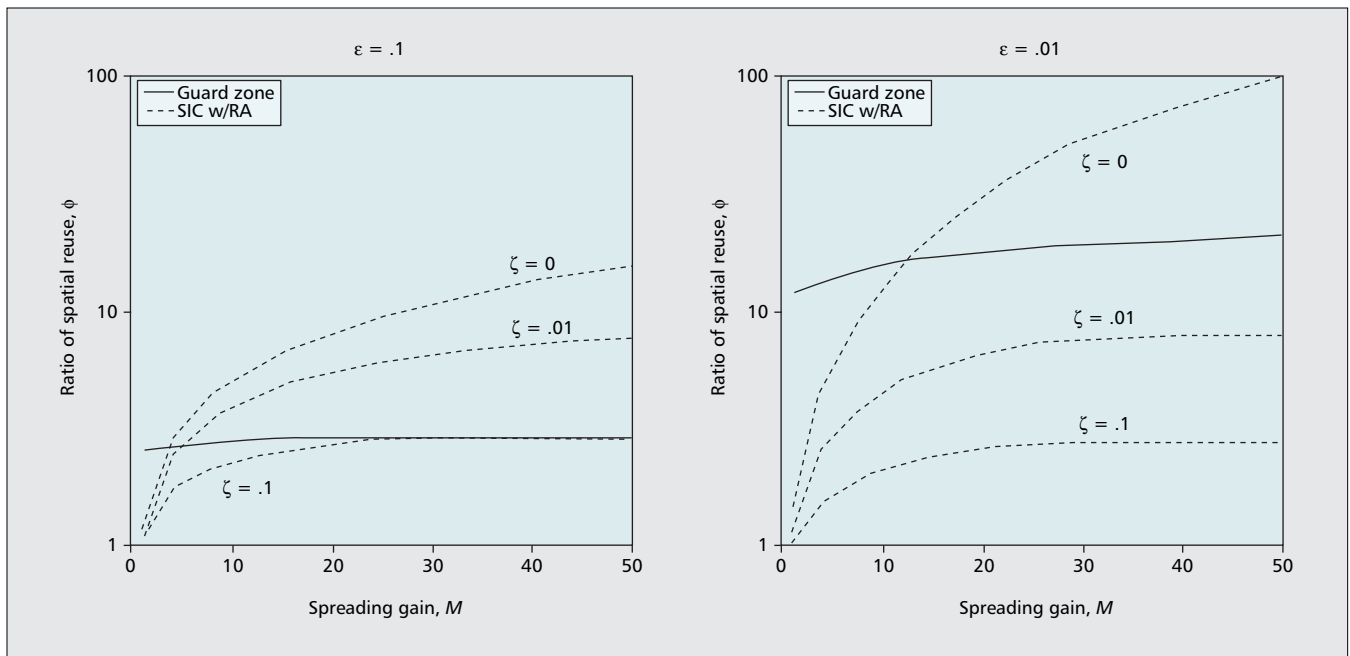
zone that should exist around each receiver; that is, the region around an active receiver that should be cleared of interfering transmissions by a higher layer protocol. There is a trade-off between protecting active receivers from outages (larger guard zone) and maximizing the number of concurrent transmissions (smaller guard zone). Scheduling in ad hoc networks is a multifaceted and rich research topic [14, references therein], but a simple way to observe the gain possible due to scheduling in ad hoc networks is to consider the optimum guard zone's effect on capacity [15]. Although guard zone scheduling is not optimal, it does require only local coordination among nodes and provides a clear connection with the network geometry concepts that are the focus of this article.

Modeling a guard zone around a receiver is conceptually similar to perfect interference cancellation; that is, preventing nodes from transmitting is like canceling them perfectly beforehand, with the important distinction that since nearby nodes are prohibited from transmitting, the number of concurrent transmissions is less than in perfect SIC. In Fig. 6 the transmission capacities are compared for two systems. The first has a scheduling algorithm that bans nodes from transmitting within a guard zone radius R_{GZ} around each desired receiver. The second system instead has SIC with varying levels of cancellation accuracy. Guard zones are as effective as 90 percent accurate SIC when the outage constraint is $\epsilon = .1$, and much more effective than SIC at low spreading gains or severe outage constraints. The former is because at low spreading gains, the dominant interferer is likely to be outside the communication range and thus impossible to cancel via SIC.

CONCLUSIONS

In the next decade, applications for autonomous wireless networks are likely to increase dramatically, and ad hoc networks — as well as their close relatives, mesh and sensor networks — will become significant components of the wireless ecosystem. Superficially, spread spectrum transmission will appear very attractive for many of these applications due to its well-known interference robustness and security features. This article is intended to deepen understanding of spread spectrum's advantages and disadvantages in the context of ad hoc networks. As we have seen, many of the design issues are quite distinct from cellular networks.

We have argued that fundamentally, interference averaging is not nearly as profitable in ad hoc networks as it is in cellular networks due to the very different geometric properties of the transmit/receive positions. Therefore, frequency hopping — interference avoidance — should generally be preferred to direct sequence spread spectrum — interference averaging. We have also noted that both FH and DS incur considerable overhead in code acquisition and synchronization, and this overhead needs to be amortized to make spread spectrum competitive. Unless new efficient schemes can be developed, this trait discourages the use of spread spectrum in ad hoc networks with high levels of mobility or bursty traffic. On the positive side, both FH and DS provide consid-



■ **Figure 6.** Guard zone performance vs. SIC as measured by the efficiency of spatial reuse ϕ , for moderate ($\epsilon = .1$) and severe ($\epsilon = .01$) outage constraints. The baseline case of no spreading, scheduling, or SIC, is $\phi = 1$.

erable flexibility to the network by allowing longer hop lengths and, in the case of DS, a reduced interference range. These aspects simplify some protocol design aspects and, perhaps most important, allow end-to-end delay and energy consumption to be reduced, perhaps substantially.

ACKNOWLEDGEMENTS

The authors extend thanks to Drs. G. de Veciana, A. Hasan, and X. Yang for their contributions to several of the ideas in this article.

REFERENCES

- [1] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless Mesh Networks: A Survey," *Comp. Networks*, vol. 47, no. 5, Mar. 2005, pp. 445–87.
- [2] M. B. Pursley, "The Role of Spread Spectrum in Packet Radio Networks," *Proc. IEEE*, vol. 75, no. 1, Jan. 1987, pp. 116–34.
- [3] T. J. Shepard, "A Channel Access Scheme for Large Dense Packet Radio Networks," *ACM SIGCOMM*, Stanford Univ., Aug. 1996, pp. 219–13.
- [4] P. Gupta and P. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Info. Theory*, vol. 46, no. 2, Mar. 2000, pp. 388–404.
- [5] S. Weber et al., "Transmission Capacity of Wireless Ad Hoc Networks with Outage Constraints," *IEEE Trans. Info. Theory*, vol. 51, no. 12, Dec. 2005, pp. 4091–4102.
- [6] A. Muqattash and M. Krunz, "CDMA-Based MAC Protocol for Wireless Ad Hoc Networks," *ACM SIGMOBILE*, June 2003, pp. 153–64.
- [7] X. Yang and G. de Veciana, "Inducing Spatial Clustering in MAC Contention for Spread Spectrum Ad Hoc Networks," *ACM MobilHoc*, May 2005.
- [8] M. Haenggi and D. Puccinelli, "Routing in Ad Hoc Networks: A Case for Long Hops," *IEEE Commun. Mag.*, vol. 43, no. 10, Oct. 2005, pp. 93–101.
- [9] M. Sikora et al., "Bandwidth- and Power-Efficient Routing in Linear Wireless Networks," *Joint Special Issue of IEEE Trans. Info. Theory and IEEE Trans. Networking*, June 2006, pp. 2624–33.
- [10] M. Haenggi, "The Impact of Power Amplifier Characteristics on Routing in Random Wireless Networks," *Proc. IEEE GLOBECOM*, San Francisco, CA, Dec. 2003, pp. 513–17.

- [11] S. Verdú, *Multuser Detection*, Cambridge Univ. Press, 1998.
- [12] J. G. Andrews, "Interference Cancellation for Cellular Systems: A Contemporary Overview," *IEEE Wireless Commun.*, vol. 12, no. 2, Apr. 2005, pp. 19–29.
- [13] S. Weber et al., "Transmission Capacity of Wireless Ad Hoc Networks with Successive Interference Cancellation," *IEEE Trans. Info. Theory*, vol. 53, no. 8, Aug. 2007, pp. 2799–2814.
- [14] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource Allocation and Cross-Layer Control in Wireless Networks," *Foundations and Trends in Networking*, vol. 1, no. 1, 2006, pp. 1–144.
- [15] A. Hasan and J. G. Andrews, "The Guard Zone in Wireless Ad Hoc Networks," *IEEE Trans. Wireless Commun.*, vol. 6, no. 3, Mar. 2007, pp. 897–906.

BIOGRAPHIES

JEFFREY ANDREWS (jandrews@ece.utexas.edu) is an assistant professor of electrical and computer engineering at the University of Texas at Austin. He received a B.S. with high distinction from Harvey Mudd College in 1995, and M.S.E.E. and Ph.D. degrees from Stanford University in 1999 and 2002. He serves as an editor for *IEEE Transactions on Wireless Communications*, and has industry experience at Qualcomm, Intel, Palm, NASA, and Microsoft. He received the NSF CAREER award in 2007 and is the Principal Investigator of an eight-university team of 13 faculty in DARPA's Information Theory for Mobile Ad Hoc Networks program.

STEVEN WEBER (sweber@ece.drexel.edu) is an assistant professor in the Department of Electrical and Computer Engineering at Drexel University, Philadelphia, Pennsylvania. He received a B.S. from Marquette University in 1995, and M.S. and Ph.D. degrees from the University of Texas at Austin in 1999 and 2003, respectively, all in electrical engineering. His research interests are focused on wireless networking, stochastic geometry, and ad hoc network capacity.

MARTIN HAENGGI (mhaenggi@nd.edu) is an associate professor of electrical engineering at the University of Notre Dame. He received Dipl. Ing. (M.Sc.) and Ph.D. degrees from the Swiss Federal Institute of Technology in Zurich (ETHZ), Switzerland, in 1995 and 1999, respectively. In 1999–2000 he spent a postdoctoral year at the University of California at Berkeley. He received an NSF CAREER award in 2005 and is a member of the Editorial Board of Elsevier's *Journal of Ad Hoc Networks*. His scientific interests include networking and wireless communications, with an emphasis on ad hoc and sensor networks.