

By P.R. Kumar

New Technological Vistas for Systems and Control:

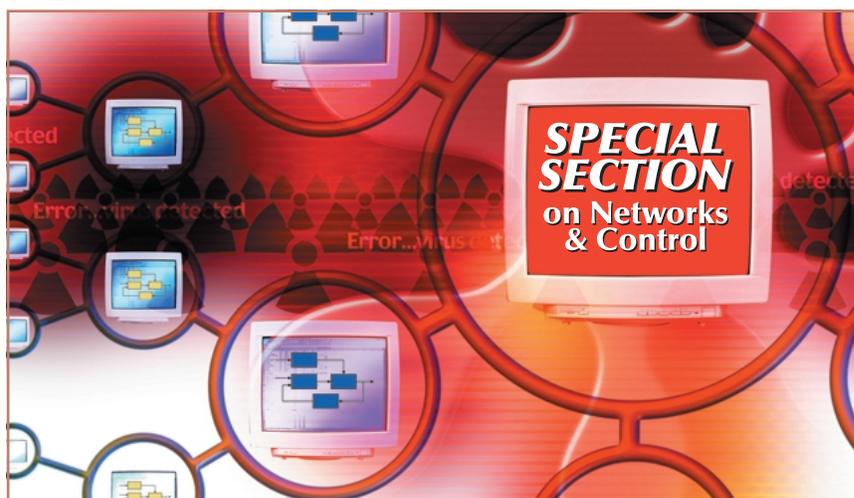
The Example of Wireless Networks

The biggest difference, though, is the control system.

—Arthur C. Clarke,

3001 The Final Odyssey. New York: Del Rey Science Fiction, 1998, p. 107.

Over the past four decades, the field of control has woven a rich tapestry of a larger systems theory, with sustained investigations into fundamental issues such as control, estimation, stability, optimality, adaptation, and decentralization. These issues are the fundamental ingredients in many new proposed technologies, which are now within our collective purview. They provide a profusion of practical examples of problems that may have sometimes abstractly engaged our attentions in the past and offer a wealth of opportunities for imaginative solutions. The opportunities are ours to seize. That is the central thesis of this article, and we use the example of wireless networking, an area of



©2000 Image 100 Ltd.

great interest in the emerging field of information technology, to illustrate it.

Research in the field of control theory has been instrumental in developing a coherent foundation for systems theory. Notions such as stabilizability, reachability, optimality, identification, adaptation, robustness, estimation, information structures, games, control with partial noisy observations, distributed control and estimation, linearity, nonlinearity, infinite dimensional systems, discrete-event systems, hybrid systems, and nonholonomic systems have been systematically explored in depth. In its search for

The author (prkumar@decision.csl.uiuc.edu) is with the University of Illinois, Department of Electrical and Computer Engineering, Coordinated Science Laboratory, 1308 West Main Street, Urbana, IL 61801, U.S.A.

deeper answers, the community as a whole has been willing to invest in learning various mathematical fields to exploit any and all tools. Complex analysis, martingale theory, heat equations, differential geometry, automata theory, functional analysis, calculus of variations, algebra, category theory, algebraic geometry, Brownian motion, and the like have all been avidly studied, and, as a result, their usage has become routine. Consequently, a rich metaphor of concepts has been created that allows us to quickly locate the central issues in a given problem formulation.

One can broadly discern two thrusts in the above research. One has been the focus on the core problem of control-loop design for linear and nonlinear differential equations. This has given us LQG control, self-tuning control, identification methodologies, H_∞ -control, l_1 -control, nonlinear control, backstepping, Kharitonov extensions, etc. Over the past decade, a profusion of software tools has been developed that has made these results widely accessible. The other thrust has been expansionary and has served to enlarge the frontiers of system theory. Topics such as learning, distributed estimation, distributed asynchronous algorithms, stochastic networks, team theory, stochastic algorithms, optimization theory, and stability have been explored. As a result, researchers with roots in control theory have made distinguished contributions to other fields such as power systems, robotics, learning theory, economics, signal processing, optimization, stochastic processes, communication networks, wireless networks, operations research, financial economics, and manufacturing systems. Our workshops and flagship conferences are routinely attended by researchers from all these disciplines. It is interesting, if only as a purely sociological phenomenon, that a field beginning with problems such as Watt's governor has gone on to make such a broad impact.

What then does the future hold? What are the vistas and opportunities, particularly in the next one or two decades? What are the current trends affecting technology in general? What lessons should we learn from the past? How can we build on our past successes and assure ourselves of an even brighter future?

Perhaps two principles could profitably be kept in mind. The first is that the control field may be well served by paying careful attention to technology and its trends. Interesting problems abound in automotive technology, semiconductor manufacturing, power systems, filter design, networking, smart materials, and so on. Moreover, future society may be radically altered by trends in nanotechnology, the deregulation of the radio spectrum, advances in materials science, biotechnology, robotics, and others. It may be beneficial for researchers to each become a domain expert in at least one technology. By learning the fundamental details of a technology, one is in a position to discover problems that may lie hidden. By starting with the basic problem, one can play a role in the very modeling of the fundamental phenomena, which opens up many possibilities

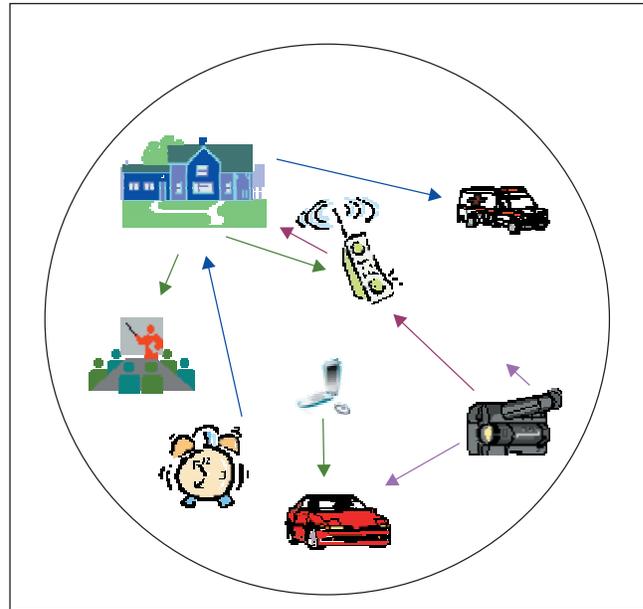


Figure 1. An ad hoc wireless network.

for design and analysis. Also, by knowing the technology, one is in a position to know where the current bottlenecks are, since advances made in bottlenecks have a greater impact than those made in nonbottleneck issues. Moreover, by starting with technology, the motivation of our research will be transparent to others. In the current utilitarian, problem-driven environment, it is important for the broader scientific and engineering community to understand the context of our work.

The second principle is that we may be well served by preserving our culture of rigorous and deep thinking and the tradition of being unafraid to invest in learning new approaches and theoretical tools. A broad and deep knowledge is invaluable in capitalizing on research opportunities. It allows one to quickly discard unworkable approaches and identify promising ones. Over the long haul, a case can be made that investment in learning does pay off.

The field of wireless networks is one fertile future vista that can serve to illustrate some of the above comments. We will see that the issues of stability, uncertainty, adaptation, information structures, and distributed control, all concepts that our discipline has researched, sometimes in abstraction, are present in practical profusion. Modeling the phenomena and identifying the research issues allows wide scope for original work. The field of wireless networks is of intense current interest and regarded as possibly revolutionary in its impact on society. There are many challenges to confront in the design of workable wireless networks and also in the development of control systems orchestrating several sensors and activators connected wirelessly. Finally, wireless networking and networked embedded devices are just two topics in the broader field of information technology, which is currently the subject of much attention.

Indeed, the central thesis of this article is this: Issues that have collectively engaged our community over the past four decades, such as optimality, estimation, stability, adaptation, robustness, uncertainty, partial observations, and distributed systems, are the fundamental ingredients in many future technologies, and by bringing to bear our tradition of rigorous analysis and design, our field has a bright future, as well as a significant role to play in the realization of several technologies. The present age is characterized by numerous attempts to build such complex systems. All these future systems can well be added to our list of core control systems.

The wireless medium is unreliable: transmissions are subject to obstacles, reflections, multipath effects, and fading.

Wireless Networks

The type of wireless networks we will describe have, at various times in the past, been called packet radio networks or multihop mobile radio networks; the current term in vogue is *ad hoc networks*.

Currently, there is much excitement about the future possibilities realizable by wireless networks. Wires, which earlier enabled the Internet revolution, have now, in many cases, simply become a barrier to proliferation. When nodes are untethered, new possibilities arise. For example, in recent decades there has been much growth in embedded devices, such as those present in alarm clocks, automobiles, cell phones, toasters, oil wells, and in fact all around us. However, they have not operated in concert, since they have been disconnected from each other. In the future, toasters may coordinate with alarm clocks, travelers may acquire maps or videos from local information warehouses, automobiles may warn each other when braking, and wireless trading may become a reality. Of course, which applications will be successful and which will not cannot be predicted. All one can say is that such technological capabilities are on the horizon. When a myriad of embedded devices are wirelessly networked, we will find ourselves confronted with legions of sensors and actuators all around us. The challenge of orchestrating all these devices into a coherent control system will loom large.

There are two roles that systems and control can play in this revolution. The first role is intrinsic; how can the principles of system and control be used to operate wireless networks? That is, systems and control researchers can enable the deployment of the wireless networks themselves. The second role is extrinsic; how can one exploit the capabilities made available by wireless networked sensors and actuators to construct next generation distributed control systems? That is, systems and control researchers will be confronted with new possibilities for pervasive control systems operat-

ing over the ether. Our focus in this article will be on the former intrinsic function of enabling wireless networks.

Consider a domain in which are located a number of nodes, as shown in Fig. 1. Each node is equipped with a wireless transmitter capable of transmitting packets to its neighbors. Packets can be relayed from node to node until they reach their final destination. Nodes may be mobile, changing their location over time. Also, nodes may switch themselves off from time to time.

The wireless medium is an unreliable one: transmissions are subject to obstacles, reflections, multipath effects, and fading, all of which affect the quality of the received signal. Moreover, transmissions can interfere with each other. All these conspire to make packet reception much more unreliable than in wired media.

The objective of operating a wireless network is to transfer packets from their sources to their destinations reliably and efficiently. Nodes must choose the power levels at which they broadcast, since that influences the range, and must time their transmissions. Nodes must also cooperate in relaying each other's packets.

Given the characteristics of the wireless medium, how does one design a system that can reliably and efficiently transport packets from sources to destinations?

The Goal: Quality of Service

Two important performance measures affect a user's experience of a wireless network, or any other communication network for that matter. One is *throughput*, which is measured in bits per second. Throughput is the rate at which bits are transferred from the source to their final destination. The second is *delay*, which is the difference in time from when a bit enters the network at its source node to when it is eventually received correctly at its destination node.

For real-time or interactive services such as voice or video, timeliness of receipt of packets is crucial. Hence, delay is an important performance measure. However, real-time voice can afford to lose some packets and still sound intelligible to the listener. Thus it may be a reasonable strategy to drop some packets if they have been excessively delayed en route or to not retransmit them if they have been corrupted en route. On the other hand, for data transfer (e.g., file transfer), delay may not be important, but accurate and complete reception may be paramount.

Other performance measures are also derived from the above performance measures. For example, the mean delay of packets, the standard deviation of the delay, or even the 99th percentile of delay may all be important indicators of performance. Another could be *fairness*. When many users are active in the network, it may be important that each receives a nearly equal throughput and/or delay. Or if some users are more important than others, for example, if they have contracted for a

greater quality of service (QoS), their throughput or delay may be weighted more heavily than others.

Although the Internet currently does not provide much in the way of guaranteeing QoS to users, future users may desire a network that has the flexibility and low cost of the Internet, but with end-to-end guarantees of the telephone network [1]. How this is to be done in the present Internet, let alone future wireless networks, is a topic of much current research. In the remainder of this article we discuss issues surrounding four problems that arise in the design of wireless networks.

The Power Control Problem

In a wireless network, a fundamental issue faced by a node on every transmission is how to choose the power level at which it should transmit. Clearly, the power level should be high enough that the intended receiver receives a signal of adequate power. It should not transmit at too high a power level, however, since that may interfere with other receivers receiving separate transmissions of interest to them. One more factor to keep in mind is that during periods of deep fading, it may be better to lose a packet and retransmit it later than to broadcast it at a very high power level. This is related to the “waterfilling” solution in information theory [2].

These issues, and others to be described later, give rise to the *power control problem*. This can be regarded as a feedback-based set-point regulation problem. The receiver can provide a feedback signal to the transmitter, allowing it to regulate its transmitted power level so that the net received power level at the receiver, or the net received signal-to-interference-plus-noise ratio at the receiver, is at a desired level. (We should mention here that every problem in wireless networks is considerably exacerbated if the links are not bidirectional.)

There is another issue to be considered, though. Several transmitter-receiver pairs may be simultaneously operating, with each having its own feedback loop (see Fig. 2). Therefore, when one transmitter increases its power level based on its own feedback, this is experienced as increased interference at another receiver receiving a different transmission, thus causing it to send a feedback signal to its own transmitter to increase its power level. Clearly, this causes coupling between loops, with each transmitter raising its power level in response to another transmitter doing so. Thus one needs to study and design schemes for power control that ensure the convergence of the power levels of all transmitters. This is a problem close to the heart of all control researchers. It can be solved, for example, by discretizing the power levels, with each transmitter raising its power level just enough to satisfy its own receiver’s desired signal-to-interference ratio (see [3]). Solutions are also available for this “physical layer” problem of maintaining a link of adequate quality when the power levels are continuously varying (see [4]-[6]).

Unfortunately, there are other complicating issues in ad hoc networks. High-power transmissions cause interference, which results in low throughput for other users. Thus the power control problem impinges on the “congestion control” problem, which is managed in the current Internet by protocols such as TCP [7]-[10]. A user’s throughput is increased until a packet is lost, at which time the throughput is drastically curtailed. This is done by giving each user a “window,” which is the number of unacknowledged packets it can have in the network, and increasing or decreasing the window. The essential point is that when the network is congested, packets are dropped at full buffers, so packet loss is taken as the appropriate “feedback” signal to regulate congestion.

In wireless networks, however, one node’s congestion can be caused by another’s usage of an excessive power level. How does one node regulate the behavior of others? How can one node use feedback across transmitter-receiver pairs? These questions require further investigation.

Another issue is that all nodes need to cooperate in maintaining system-wide network connectivity. It is shown in [11] that if n nodes are randomly (uniformly and independently) located in a disk of area A m², and each node chooses the same range $r(n)$, then the network is asymptotically connected with probability one as $n \rightarrow \infty$ if and only if

$$r(n) = \sqrt{\frac{\log n + \gamma(n)}{\pi n}}$$

with $\gamma(n) \rightarrow +\infty$. Thus nodes will have to cooperate with each other in choosing their individual power levels so as to ensure connectivity of the resulting network.

Finally, it is shown in [12] that under certain models, the system-wide transport capacity of the wireless network is optimized when every “hop” covers a very short distance (also see the later discussion on how much traffic wireless networks can carry). Nodes should therefore relay packets over very short distances to nearby nodes, allowing them to transmit at low power. (As noted above, however, care must be taken to prevent disconnectivity of the network.)

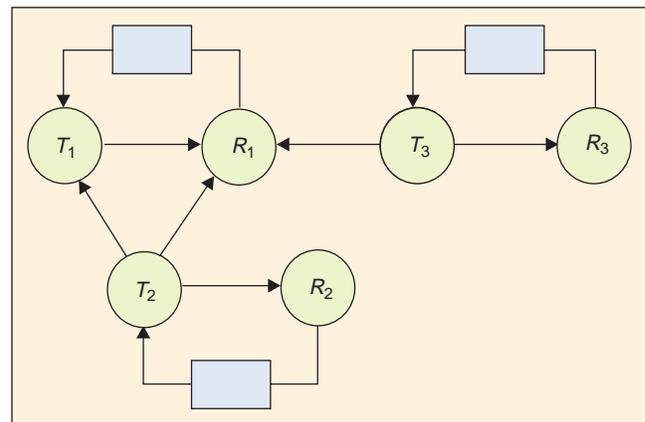


Figure 2. Coupling between power control loops of transmitter-receiver pairs.

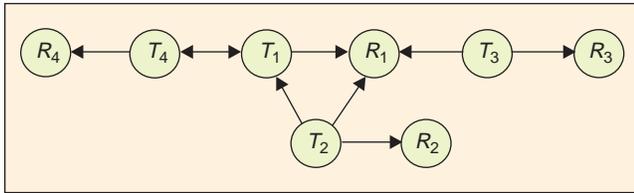


Figure 3. Interference and hidden nodes.

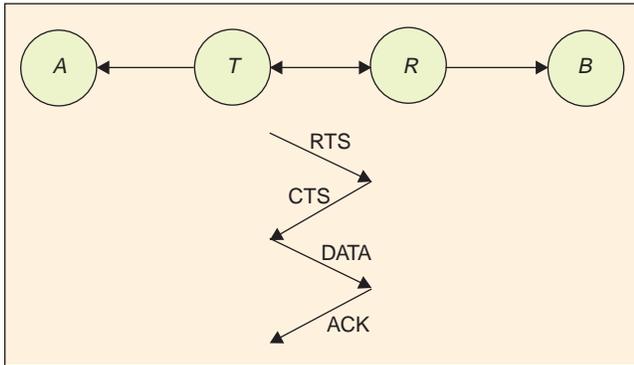


Figure 4. The nodes involved in an RTS-CTS-Data-ACK handshake in IEEE 802.11.

This will entail a packet being relayed repeatedly, traversing several hops before reaching its destination. This strategy is preferable to broadcasting at a very high power level and reaching the remote destination in just one hop, since it increases the throughput capacity that it can provide to all nodes. It does entail increased delay since packets suffer delay in traversing many hops, though of course that may be the price one has to pay to provide high throughput to all users. Thus power control affects the capacity of the system and also impinges on the routing problem considered later.

As can be seen from the above discussion, power control cuts across many layers in the design hierarchy of wireless networks and hence requires proper conceptualization.

The Medium Access Control Problem

Several problems in wireless networks arise from the fundamental shared nature of the wireless medium. One is the power control problem described previously. Another is the medium access control problem described here.

For a receiver to intelligibly receive a packet, other nearby transmitters should refrain from broadcasting. Otherwise, the phenomenon of “collisions” occurs where packets destructively interfere with each other.

Consider the scenario shown in Fig. 3. Assume that T_1 , T_2 , and T_3 are all within range of the reception by R_1 . Then, for R_1 to successfully receive a packet from T_1 , both T_2 and T_3 have to remain silent.

Note that “carrier sensing” does not guarantee that conflicts can be avoided. For example, if T_1 can hear T_2 , then it can sense the carrier of T_2 and refrain from transmitting at a time when T_2 is transmitting. This avoids collisions with T_2 . However, such a strategy does not work in avoiding collisions

with all potential interferers. Suppose T_3 can be heard by R_1 but cannot be heard by T_1 . Such a terminal T_3 is said to be a “hidden” terminal [14]. Even though T_3 can cause collisions at R_1 , T_1 cannot detect when T_3 is transmitting and thus when collisions will occur. Finally, carrier sensing can even result in inefficiency. T_1 need not refrain from transmitting to R_1 just because it hears T_4 's carrier, since T_4 's transmission does not cause interference at R_1 .

This gives rise to the *medium access control problem*. How should nodes schedule their transmissions to have their packets received intelligibly by their intended receivers?

This problem has its genesis in ALOHA, a packet radio communication system set up in the 1970s to link the Hawaiian islands [15]. (Apparently, the goal of communication at that time was only remote login; e-mail was not envisioned!) The ALOHA protocol consisted simply of transmitting whenever there was a packet to send. If there was a collision, it was assumed that all nodes could detect it, which is the case if all transmissions can be heard by all nodes or if all transmissions can be heard by a common satellite. All nodes involved in the collision would then “back off” for a random time and retry their transmissions. Note that it is necessary for the retry interval to be random; otherwise, if each node involved in the collision waits exactly the same time before retrying, then all nodes are destined to collide again. Calculations show that the maximum throughput of such a protocol over a common channel is only $1/2e$. This can be doubled by “slotting” time into intervals and synchronizing all transmissions to fit into these time slots. One adverse property of the ALOHA protocol is that if a very large number of users collide, there are likely to be collisions in the future as nodes retry. This causes the throughput of successful transmissions to drop, leading to even more backlogged transmissions as more packets arrive at the system. Put simply, ALOHA is unstable when there is an unlimited number of users. Variants of ALOHA exist today, the most notable being the Ethernet protocol, where of course the number of users is bounded.

Ad hoc networks differ from the one common channel or satellite channel scenario, however, since not all nodes can hear all transmissions. Indeed, as we have seen, there can exist “hidden” transmitters of whose existence one is not even aware. Our goal is to spatially reuse the radio frequency spectrum; that is, by restricting the range of a transmission, one can limit the interference of one's transmission, thus allowing a distant transmitter-receiver pair to carry on a conversation on the same frequency at the same time. Therefore, transmissions need to be spatiotemporally scheduled. One scheme that has been proposed is embodied in the IEEE 802.11 standard [16] (see also [17] and [18]). This protocol employs reservation packets to reserve the channel locally in space for data packets.

When a node T has a data packet to send to a node R , it first sends out a “request to send” packet, called RTS (see Fig. 4). (Here we are assuming that the node T is not cur-

rently under an order to remain silent, as described in the sequel.) This transmission can be heard by all the neighbors of T , including R . Suppose, as in Fig. 4, that T has one additional neighbor, A , and R has one additional neighbor, B . If A is not currently in range of any other transmissions, then it can hear T 's RTS. It should then refrain from transmitting for a while. Similarly, if R is not currently in range of any other transmissions, then it too can hear T 's RTS. Assuming that it is not currently under an order to remain silent, it sends a "clear to send" (CTS) packet to T . This CTS packet will be heard by B , assuming that it is not in range of any other transmission, and it should then refrain from transmitting for a while. Assume this to be the case. The CTS from R is heard by T , since A was previously silenced and therefore it does not cause interference by transmitting. T then sends its data packet to R , which receives it successfully since its neighbor B has been silenced. R then sends back an "acknowledgment" (ACK) packet to T . At that point, B is released from silence. Node A also is released from silence when it hears the end of T 's data packet, after an obligatory pause to allow node T to receive the subsequent ACK from node R . In the event of a failure for any reason (e.g., not receiving a CTS), node T simply restarts the whole process after a random time.

This handshake needs to be done for each and every data packet on each of its hops.

Several issues must be considered in evaluating this protocol: What is the overhead experienced by the RTS-CTS reservation process? How efficient is this protocol? It should be noted that defining "efficiency" is not as easy as it may seem, since it depends on the geographical location of all nodes. Indeed, due to spatial re-use of frequency, efficiency at one location cannot be defined independently of other locations. Moreover, efficiency cannot be separated from "fairness." How equitably does this protocol allocate throughput when there are several contending sources? How does this fairness depend on geometry (i.e., the location of contenders and other nodes)? How can one provide more throughput to users requesting more? What is the delay per hop? How does the protocol function in a multihop scenario where a receiver has to transmit (i.e., relay) every packet it receives? How is the protocol influenced by channel errors? Note that if a node does not successfully receive a reservation packet (e.g., RTS or CTS), then it is liable to transmit at an inopportune time when it is supposed to have been silent, thus causing a collision. The answers to many of these questions are frequently based only on simulation.

It is worth noting that the central problem here is one of achieving *distributed coordination*. To communicate, the nodes need to coordinate their transmissions; however, this coordination can itself take place only over the communication medium. So, to coordinate, the nodes need to communicate. Thus we have a vicious circle: Communication needs coordination, which itself needs communication. Can we communicate somewhat more with perhaps

fewer coordination packets? The following new protocol is a suggestion in this regard but requires further evaluation and development.

The SEEDEX Protocol

The key idea behind the proposed new SEEDEX protocol is for each node to publish a schedule of its intentions. Once all concerned nodes know of each other's schedules, they can find opportunities for transmission of their own packets.

The concept raises several questions, the first of which is what sort of schedule to publish. Our suggestion is to publish a pseudorandom schedule (i.e., a stochastic process). A second question is how to publish an entire stochastic process. The key idea we will employ is to publish the initial *seed* for a pseudorandom number generator, since that will determine the entire future of the random schedule. A third question is how to exploit this schedule to find transmission opportunities for individual transmitter-receiver pairs. Here we will simply choose the transmission times as the times when a node is certain that its intended receiver is sure to be listening, and then transmit with a probability that reduces the likelihood of a collision. Finally, we design a mechanism to acknowledge packets so that transmitters become aware of packets that have undergone collisions and can therefore retransmit them.

We suppose that time has been divided into slots of equal length. In every slot, a node is in either a "listen" mode denoted by the state L or a "possibly send" mode denoted by the state S . Each node i chooses to be in state S with probability $p(i)$ or in state L with probability $(1 - p(i))$, independently in each slot. Thus the dynamics of the state of each node can be modeled as a discrete-time, two-state Markov

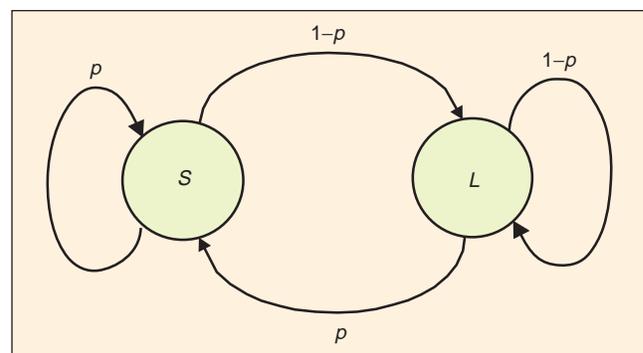


Figure 5. The state transition diagram of a node.

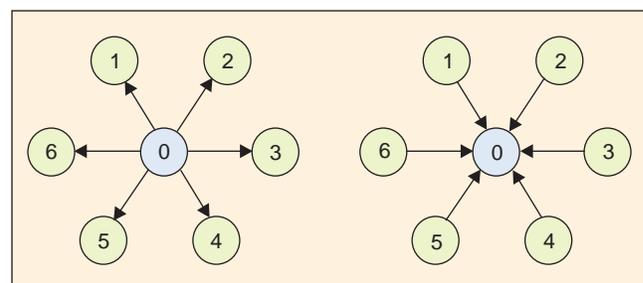


Figure 6. The two phases of the seed exchange procedure.

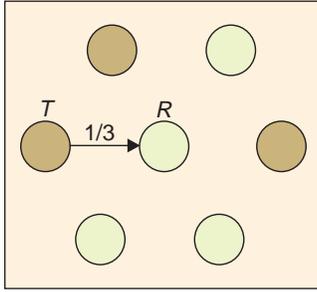


Figure 7. The transmission probability of the node T is $1/3$ since its intended receiver R has two other nodes in state S . All nodes in state S are shaded.

chain, with transition probabilities as shown in Fig. 5. Each node implements this Markov chain by employing a pseudorandom number generator, which is a simple linear congruence relationship whose output produces numbers that appear to be random. If the initial state of the pseudorandom generator, known as the *seed*, is specified, then the entire sequence of pseudorandom numbers is predetermined. Thus, once other nodes are made aware of node f 's seed, they know what state node i is in each subsequent slot. Therefore, to publish its entire future schedule, node i only needs to let others know its initial seed. Note that a node's packet can collide only with the packets of nodes in its two-hop neighborhood (recall the notion of a "hidden" node discussed earlier), and hence only such nodes need to communicate their seeds with each other.

This is done using the two-phase seed exchange procedure illustrated in Fig. 6, from which SEEDEX derives its name. In the first phase, every node transmits its seed to each of its neighbors. This can be done by simple broadcast. In the second phase, each node broadcasts the seeds of all its neighbors to all its neighbors. Thus, at the end of the two phases, every node finds out the seeds of all its neighbors, as well as the seeds of the neighbors of all its neighbors. In other words, each node is made aware of the seeds of all nodes in a two-hop neighborhood around itself. In the first phase, shown at the left in Fig. 6, node 0 broadcasts its seed to its six neighbors: 1, 2, 3, 4, 5, and 6. Similarly, in the first phase, other nodes also broadcast their own seeds to their neighbors. Thus, at the end of the first phase, all nodes know the seeds of their one-hop neighbors. In the second phase,

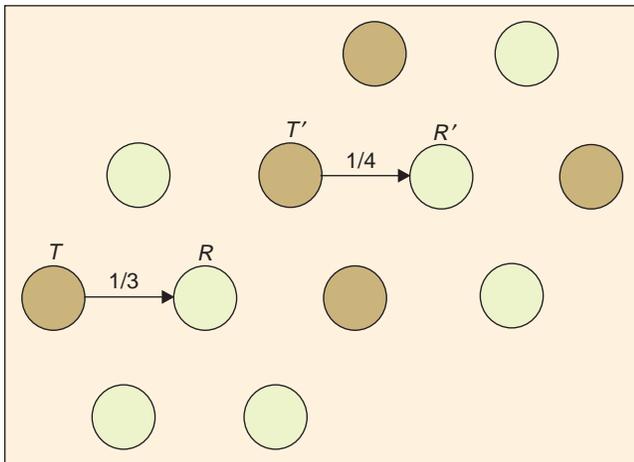


Figure 8. Node T transmits with probability $1/3$, whereas node T' transmits with probability $1/4$.

shown at the right in Fig. 6, nodes 1, 2, 3, 4, 5, and 6 broadcast the list of seeds of their own neighbors to node 0. Thus, at the end of this phase, node 0 knows the seeds of all its two-hop neighbors.

Note that each phase will take several time slots. The reason is that all nodes cannot broadcast simultaneously. They may do so randomly, á la ALOHA. Also it is not necessary to broadcast the initial seed itself; the nodes can broadcast the current seed (i.e., the current state of the pseudorandom number generator).

The next issue is to choose transmission opportunities. Suppose node T has a packet to send to node R . It waits for a slot in which it is in state S and node R is in state L , guaranteeing that node R is in a listen mode. There is another complication, however. There may be other neighbors of node R that are also in state S and thus also liable to transmit, which would result in a collision. Due to the seed exchange procedure, however, node T knows precisely how many such nodes there are; let this number be n . It then transmits on that slot with probability $1/(1+n)$. (This can be replaced by $\text{Min}[\alpha/(1+n), 1]$, where $\alpha > 0$ is some parameter to be optimized.) The idea here is to reduce the probability of transmission if there are many possible interferers and to increase it if there are few or none. This is illustrated in Fig. 7 where node T transmits with probability $1/3$. Note that if each of the other n neighbors of R actually had a packet to send to R , they too would each transmit with probability $1/(1+n)$, which would lead to an expected value of only one transmission, which is what we aim for to avoid a collision. In that case, the throughput μ_T in packets per slot that node T obtains for its transmissions to R is

$$\mu_T = p(1-p_R) \sum_{i=0}^{F-1} \binom{F-1}{i} p^i (1-p)^{F-1-i} \left(\frac{1}{1+i} \right) \left(\frac{i}{1+i} \right)^i \quad (1)$$

assuming that all nodes other than the receiver have $p_i \equiv p$ and that F is the number of transmitting neighbors of R (i.e., "flows" through R). For $F=3$ and $p_R = Fp$, this is maximized at $p^* = 0.118$, giving a throughput of about $\mu_T = 0.049$ packets per slot.

Generally though, the other neighbors of node R that are in state S may not have any packets to send, or if they do, they may want to send them not to R but to another of their neighbors. In that case, such a node will transmit with a different probability $1/(1+n')$, where n' is the number of neighbors of its recipient R' that are in state S , assuming of course that R' is in state L . Hence the above calculation is only an approximation, as illustrated in Fig. 8. To assess the "efficiency" of this scheme is therefore more complicated.

To roughly evaluate how good the performance is, consider the scenario shown in Fig. 9, where there is one flow of packets from node 0 to node 1 to node 2 to node 3 to node 4. Suppose the flow achieves μ packets per slot. Note that when node 1 is receiving a packet (from node 0), node 2 cannot transmit. Similarly, when node 1 is transmitting to node

2, node 2 cannot transmit. Thus, out of three slots, node 2 can transmit in only one of them. Hence μ cannot exceed $1/3$. If we normalize by this upper bound of $1/3$, we obtain a rough measure of “efficiency,” $\mu/(1/3) = 3\mu$ (it is actually a lower bound on efficiency). Note that there is one more constraint, which we have ignored; node 2 cannot receive a packet while node 3 is transmitting to node 4.

To extend this to a situation where there is more than one

How does one design a system that can reliably and efficiently transport packets from sources to destinations?

flow through a node is even more complicated, since there is some simultaneity in transmission opportunities which could be exploited. Consider the situation shown in Fig. 10, where there are three flows through node 0. Nodes 3 and 1 can receive at the same time while node 0 is silent. Nevertheless we will take $3 \times$ (sum of throughputs through a node) as a measure of the efficiency of the scheme. By this measure, from (1), with $p^* = 0.118$, we see that the efficiency is about 44%.

Let us now turn to the issue of QoS. Consider a scenario where a node 0 has F transmitters, $1, 2, \dots, F$, that have packets to send to it. Let p_i be the probability that node i is in state S in a slot. Denote by μ_i the service rate in packets per slot that node i obtains. Then, similar to (1), we have

$$\begin{aligned} \mu_1 &= p_1(1-p_0) \sum_{\substack{0 \leq k_2 \leq 1 \\ 0 \leq k_3 \leq 1 \\ \dots \\ 0 \leq k_F \leq 1}} \left(\prod_{i=2}^F p_i^{k_i} (1-p_i)^{1-k_i} \right) \\ &\quad \times \left(\frac{1}{1 + \sum_{i=2}^F k_i} \right) \left(1 - \frac{1}{1 + \sum_{i=2}^F k_i} \right)^{\sum_{i=2}^F k_i} \\ &= p_1(1-p_0) E \left[\left(\frac{1}{1 + \sum_{i=2}^F k_i} \right) \left(\frac{\sum_{i=2}^F k_i}{1 + \sum_{i=2}^F k_i} \right)^{\sum_{i=2}^F k_i} \right] \end{aligned}$$

(where the k_i s are independent Bernoulli random variables with probability p_i of being 1 and $(1-p_i)$ of being 0)

$$\begin{aligned} &= p_1(1-p_0) E \left[\frac{1}{1+x} \left(\frac{x}{1+x} \right)^x \right] \quad (\text{where } x := \sum_{i=2}^F k_i) \\ &\geq p_1(1-p_0) \frac{1}{1+E(x)} \left(\frac{E(x)}{1+E(x)} \right)^{E(x)} \end{aligned}$$

(by Jensen’s inequality, since $[1/(1+x)][x/(1+x)]^x$ is convex)

$$\begin{aligned} &= p_1(1-p_0) \frac{1}{1 + \sum_{i=2}^F p_i} \left(\frac{\sum_{i=2}^F p_i}{1 + \sum_{i=2}^F p_i} \right)^{\sum_{i=2}^F p_i} \\ (\text{since } E(x) &= \sum_{i=2}^F p_i) \\ &\geq \frac{p_1(1-p_0)}{14 + e \sum_{i=2}^F p_i} \end{aligned}$$

(since $[1/(1+y)][y/(1+y)]^y \geq [1/(14+ey)]$).
Similarly, for $i=2, \dots, F$, we can obtain

$$\mu_i \geq \frac{p_i(1-p_0)}{14 + e \sum_{j=1, j \neq i}^F p_j}$$

Note that this provides a relationship between the parameter setting p_i and the service rate μ_i that is obtained for a flow. Also

$$3 \sum_{i=1}^F \mu_i \geq \frac{3(1-p_0) \sum_{i=1}^F p_i}{14 + e \sum_{j=1}^F p_j}$$

As argued earlier, we take the left-hand side as a measure of efficiency.

Assume now that

$$0 \leq p_i \leq \bar{p} < 1 \quad \text{for all } i \quad (2)$$

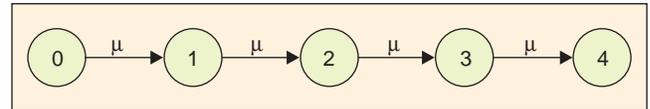


Figure 9. A flow of packets.

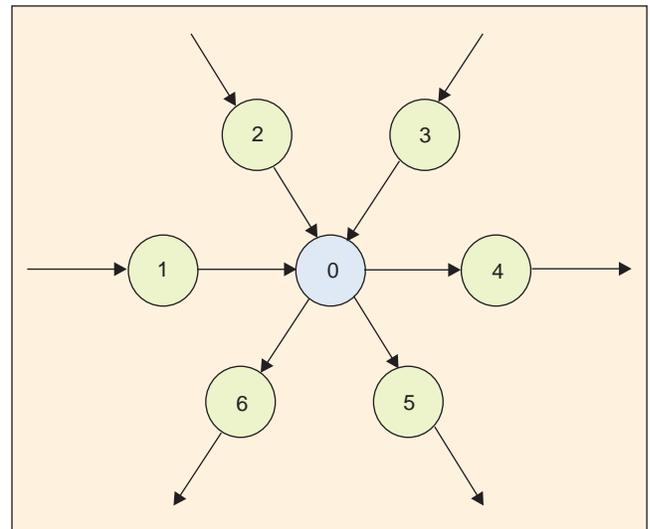


Figure 10. Many flows through a neighborhood.

and

$$\sum_{i=1}^F p_i \geq \underline{P}. \quad (3)$$

Then

$$3 \sum_{i=1}^F \mu_i \geq \frac{3(1-\bar{p})\underline{P}}{14 + e\underline{P}}$$

as well as

$$\frac{\mu_i}{p_i} \geq \frac{1-\bar{p}}{14 + e(F-1)\bar{p}}. \quad (4)$$

This latter lower bound can be taken as a QoS guarantee on the guaranteed throughput, when the guidelines (2), (3) are used to allocate the p_i s at each node. The numbers shown in Table 1 for $F = 2, 3, 4, 5, 6$ are lower bounds only, and the performance will generally be better.

The next issue that naturally arises is how much μ_i to “provide” to a requesting flow and how to allocate its p_i . One scheme that suggests itself is that, when setting up a flow, a leader packet traverses the route gathering information on how much service rate is available at each node. On the forward path, it finds out where the bottleneck is, and on the reverse path, it requests the minimum service rate that can be provided at the nodes it has transited and obtains an appropriate p_i at each node. Several phenomena need to be studied in this regard. Altering any p_i also alters the service rate obtained by other already existing flows. This causes a reverberating distributed adjustment mechanism whose equilibration is of interest.

It is useful to note that regulating the service rate obtained by a flow at each node it traverses opens up the possibility of congestion control on a hop-by-hop basis. If λ_i is the throughput of the flow and μ_i is the service rate, then approximation by an $M/M/1$ queue would show that the number of backlogged packets has the probability distribution $\pi(n) = (1-\rho_i)\rho_i^n$ with a mean of $\rho_i/(1-\rho_i)$ and a standard deviation $\sqrt{\rho_i/(1-\rho_i)}$, where $\rho_i = \lambda_i/\mu_i$. (An $M/M/1$ queue is a queue to which customers arrive as a Poisson process, and which is served by one server, with service times being independent and exponentially distributed.) This suggests strategies for cutting the flow when the buffer length is unduly large and increasing it otherwise, which can be incorporated in a window-based mechanism such as TCP.

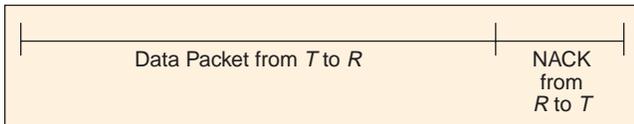


Figure 11. Structure of a slot.

Table 1. Guaranteed service rate and efficiency for $\bar{p} = 0.3, \underline{P} = 0.6$.		
Number of Flows F	Lower Bound on $\frac{\mu_i}{p_i}$	Lower bound on $3 \sum_{i=1}^F \mu_i$
2	0.32	0.42
3	0.23	0.42
4	0.18	0.42
5	0.15	0.42
6	0.13	0.42

To address the issue of how acknowledgments can be obtained, consider the structure of a slot as shown in Fig. 11. The bulk of the slot is reserved for the data from T to R (as well as the header bits); however, the tail end of the slot is reserved for a “negative acknowledgment,” or NACK. If packets numbered $1, 2, 3, \dots, k-1, k+1, \dots$ have been received, but not k , then the receiver requests packet k . Note that the receiver does this persistently until it obtains packet k , at which time it may switch to requesting the next missing packet. No protection against collisions is provided to this reverse channel; however, since it is persistent, it will eventually be heard. One also needs to calculate the throughput on the reverse channel to show that missing packets do not become the bottleneck for the forward link.

Several other properties of the proposed protocol may be noted. One is that a node with only a packet or two still gets opportunities to transmit, even without reserving a flow. Such an opportunity occurs whenever all neighbors of the intended receiver are in state L , as well as the receiver itself, guaranteeing an unhindered path for its packet. One other point is that the protocol is relatively unaffected by channel errors, since these are simply treated as collisions that are retransmitted following NACKs. Finally, we note that one can consider schemes which employ SEEDEx in an RTS-CTS handshake scheme with the goal of improving throughput.

The Routing Problem

In an ad hoc network, the address of a node is not indicative of its location or how to reach it. The address is simply a name. This gives rise to the routing problem [19], where the goal is to determine the route to be followed by packets from their sources to given destinations.

Suppose one wishes to find the path with the minimum number of hops from a given node 0 to a particular node i . This can be viewed as a dynamic programming problem. Let d_{ji} be the shortest path from node j to node i . For each node j , let N_j denote the set of one-hop neighbors of j . Assume, for simplicity, that all links are bidirectional and that the network is connected. The d_{ji} uniquely satisfy the dynamic programming equations:

$$d_{ii} = 0,$$

$$d_{ji} = 1 + \text{Min}_{k \in N_j} d_{ki} \text{ for } j \neq i.$$

This can be solved using the Bellman-Ford algorithm [20]. Let $d_{1k}^{(h)}$ be the shortest path from 1 to k with the constraint that the total number of hops is less than h . $d_{1k}^{(h)}$ is defined as $+\infty$ if there is no path from 1 to k with fewer than h hops. Define

$$d_{11}^{(h)} := 0 \text{ for all } h$$

Initially, let

$$d_{1k}^{(0)} = +\infty \text{ for all } k \neq 1.$$

Then, for $h=1,2,\dots$, let

$$d_{1k}^{(h+1)} = \text{Min}_{j \in N_k} [1 + d_{1j}^{(h)}]$$

for all $k \neq 1$.

Another solution is to use Dijkstra's algorithm. The procedure employed is to find the shortest paths to nodes in the order of increasing path length. Let X be the set of nodes to which we have obtained the shortest paths from node 1. Set

$$X = \{1\}, \quad d_{11} = 0, d_{1j} = 1 \text{ for all } j \in N_1, j \neq 1,$$

$$\text{and } d_{1j} = +\infty \text{ for all } j \notin N_1 \cup \{1\}.$$

At each step we then perform the following operations (X^c refers to the complement of X):

- i) Find $i^* \in X^c$ such that

$$d_{1i^*} = \text{Min}_{j \in X^c} d_{1j}.$$

- ii) Set $X = X \cup \{i^*\}$.
- iii) Set $d_{1j} = \text{Min}[d_{1j}, d_{1i^*} + 1]$ for all $j \in N_{i^*} \cap X^c$.

Distributed versions of these algorithms also exist. For example, in the distributed asynchronous Bellman-Ford algorithm [21], a node can broadcast its current estimate of its distance to its neighbors. Each neighbor then adds one to this distance (the additional one is for the added hop) and compares this quantity with its own current estimate of its shortest distance. If the new distance is less, then the node adopts it and also notes the node through which packets should be sent.

In a dynamic context, when nodes are mobile and links are volatile, problems can arise. One such problem is the "counting to infinity problem." Consider the scenario shown in Fig. 12. In the beginning, there is a path from node A to node C via node B . The distances are $\hat{d}_{BC} = 1$ and $\hat{d}_{AC} = 2$, and both nodes A and B are aware of this. We use a "hat" to signify that these are merely estimates of the distances as perceived by the nodes. Now suppose the link from B to C is broken. Node B then notes that A had a path to C of two

hops, and since A is its neighbor, it updates its estimated distance to C as $\hat{d}_{BC} = 3$. When this is passed on to A , it revises its distance to C as $\hat{d}_{AC} = 4$, and so on.

This illustrates the problem of maintaining routes that are loop free. Other issues of importance are to reduce the overhead in the amount of information transported just to maintain routes and to reduce the time that routing algorithms take to reconverge to a solution when nodes move or links fail.

Routing protocols are broadly classified as based on either "link state" or "distance vector." In link state protocols, each node broadcasts its complete view of the network to all nodes. Hence each node can form its own picture of the complete network. In distance vector protocols, each node only broadcasts its view of its distances to other nodes (as in the distributed Bellman-Ford algorithm, for example). An early example of a distance vector protocol is [22], whereas a more recent one that employs sequence numbers to avoid the formation of loops is [23].

Several protocols have been proposed for use in ad hoc networks. In the Dynamic Source Routing protocol [24], the sender explicitly lists the route of the packet in the header of each packet. When a route is needed and is not available in cache, a route discovery procedure is used. A route request packet is repeatedly rebroadcast from node to node until it reaches its final destination, where the reversed route can be used to intimate to the source the sequence of nodes followed by the route discovery packet. Several modifications are made to maintain loop-free routes, cache routes, etc. In the Adhoc On-Demand Distance Vector Routing protocol [25], an attempt is made to preserve the advantages of distance vector routing while eliminating the overhead of global periodic routing advertisements. Several modifications are made to repair broken paths and to maintain loop-free routes. In the Temporally Ordered Routing Algorithm [26], shortest path routing is eschewed. The goal is to construct and maintain a set of directed acyclic graphs rooted at the destinations. Such graphs can provide a multiplicity of routes to a destination.

An idea suggested in the System and Traffic Adaptive Routing Algorithm [27] is to use an adaptation algorithm to adapt routes to the network delays. Let $D_{sk}^d(t)$ denote an estimate made at time t of the mean delay from source s to desti-

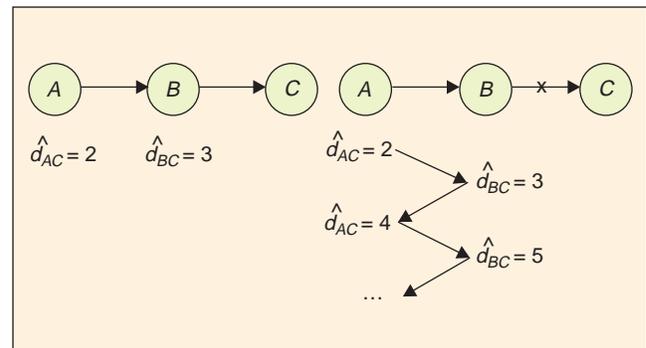


Figure 12. The counting to infinity problem.

nation d , if the packet is sent on its first hop to node k . In this algorithm, packets at source s destined for node d are sent on their first hop to node k with probability p_{sk}^d . These probabilities are adapted according to the following adaptation law:

Two important performance measures affect a user's experience of a wireless network: throughput and delay.

$$p_{sk}^d(t) = p_{sk}^d(t-1) + \alpha(t)[D_s^d(t-1) - D_{sk}^d(t-1)],$$

with some modification to keep them nonnegative and summing to one, where $D_s^d(t-1)$ is an estimate of the average delay:

$$D_s^d(t-1) = \sum_{k \in N_s} p_{sk}^d(t-1) D_{sk}^d(t-1).$$

The parameter $\alpha(t)$ is an adaptation step-size parameter. The adaptation update reduces the probability of forwarding a packet via node k if the estimated delay of the route via k is larger than the average estimated delay and increases it otherwise.

Thus, in equilibrium, the probabilities will be such that the mean delay via every utilized relay node k is the same as the average delay, and all the routes via the unutilized relay nodes k have a greater potential delay.

Hence: Every path between nodes s and d that has a positive fraction of traffic traversing it has the same mean delay. All unutilized paths have a larger mean delay.

This type of equilibrium is apparently well known in transportation; it is called a Wardrop equilibrium [28].

Clearly, the central problem is to design an efficient adaptive distributed routing algorithm. It must be adaptive due to the volatile and mobile nature of the network, and it needs to be distributed, since the nodes are neither spatially collocated nor informationally centralized. One needs to understand the time constants of adaptation and how

they compare with the time constants of mobility and also the tradeoff between the value of maintaining routes and the overhead associated with doing so.

How Much Traffic Can Wireless Networks Carry?

Given the inherent shared nature of the wireless medium, what is the ultimate traffic-carrying capacity of wireless networks? The answer will provide a design goal to strive for.

To answer this question, one needs to model the nature of interference to receivers, or its converse: When is a packet successfully received by a receiver? Consider a model in which there are n nodes in a disk of area A m², with each node capable of transmitting at W bits per second. Nodes can choose the range of each transmission. We will suppose that a transmission of range r creates a wireless footprint in a disk of radius $(1 + \Delta)r$ centered around the transmitter, within which no receiver can successfully receive any other transmission, as shown in Fig. 13. (For more details, as well as other models centered around the receiver, or based on the signal-to-interference ratio, see [12] and [13].)

The performance measure we will analyze is the aggregate *bit-meters/second* (bit-m/s) that the network as a whole can transport. We call this the *transport capacity* of the network. When 1 bit has been transported 1 m, we will say that the network has pumped 1 bit-m. Counting all such transactions occurring in the network, how many bit-meters per second can the network pump?

For simplicity of exposition, suppose that time is slotted into slots of length τ , and transmission is quantized into packets of length τs that fit into slots. Consider a long time interval of length T , and suppose that in this time interval a total of B bits have been transported from their sources to their destinations, in a multihop manner. Let us suppose that the average distance traveled by a bit is L m; thus the amount pumped by the entire network is BL/T bit-m/s.

We will show that

$$\frac{BL}{T} \leq \sqrt{\frac{8}{\pi}} \cdot \frac{W}{\Delta} \sqrt{An} \text{ bit-m/s.} \quad (5)$$

This provides an upper bound on the transport capacity of the network. On the right-hand side, the linear scaling in W is obvious. If everyone doubles the capacity of their modems, the number of bits is also doubled. Similarly, the square root scaling in area A is also obvious. If the area is doubled, distances are scaled by the square root of two.

Of most interest is the square root scaling in the number of nodes n , which has the following implication. If the transport capacity is divided equitably among all the nodes, then

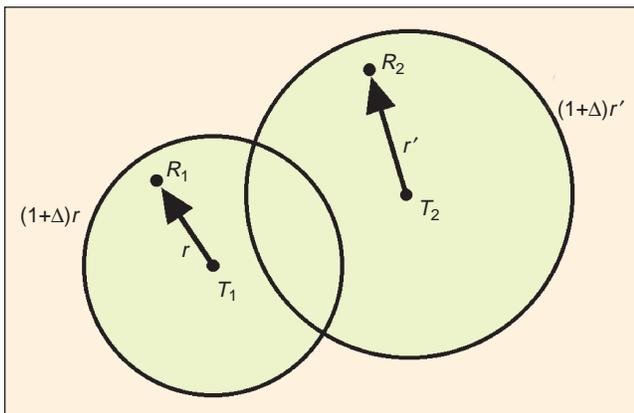


Figure 13. The interference footprint of transmissions.

each node individually obtains only $O(1/\sqrt{n})$ bit-m/s. Thus, as the number of nodes in the network increases, there is a square root drop-off in what each node obtains.

Hence networks should be designed either with few nodes (i.e., small n) or to support mainly nearest neighbor communications (i.e., small distances). Note that the distance to the nearest neighbor is of order $O(1/\sqrt{n})$ m when there are n nodes in a domain.

We will now demonstrate the square root scaling, which follows from some fundamental constraints. From Fig. 13, it is obvious that

$$\begin{aligned} |R_2 - T_1| &\geq (1 + \Delta)|R_1 - T_1|, \\ |R_1 - T_2| &\geq (1 + \Delta)|R_2 - T_2|. \end{aligned}$$

Hence, by the triangle inequality,

$$\begin{aligned} |R_2 - R_1| + |R_1 - T_1| &\geq (1 + \Delta)|R_1 - T_1|, \\ |R_1 - R_2| + |R_2 - T_2| &\geq (1 + \Delta)|R_2 - T_2|. \end{aligned}$$

Thus

$$\begin{aligned} |R_2 - R_1| &\geq \Delta|R_1 - T_1|, \\ |R_1 - R_2| &\geq \Delta|R_2 - T_2|. \end{aligned}$$

Hence

$$|R_2 - R_1| \geq \frac{\Delta}{2} (|R_1 - T_1| + |R_2 - T_2|).$$

This shows that disks of radius $\Delta/2$ (range) centered around receivers are disjoint. Thus each transmission of range r consumes an area of $(\pi r^2 / 4 \Delta^2)$. If a receiver is close to the boundary of the domain, some of the consumed area can lie outside the domain, as shown in Fig. 14; however, at least a quarter, or $(\pi r^2 / 16 \Delta^2)$, lies within the domain.

Now we simply list some of the constraints on the operation of the network. First, we need to do some bookkeeping. Suppose that the b th bit makes $h(b)$ hops on its way from its source to its destination, with the h th hop moving a distance of $r(b, h)$ m. Note that each hop of a packet occupies one slot and carries τW bits.

Clearly

$$\sum_{b=1}^B \sum_{h=1}^{h(b)} r(b, h) \geq BL \quad (6)$$

since the total distance traveled by a bit can be no less than the shortest path from its source to its destination.

If we count all the bits undergoing hops, we obtain the bound

$$H := \sum_{b=1}^B h(b) \geq \frac{WTn}{2} \quad (7)$$

since at any given time at most $n/2$ nodes can be transmitting, each at rate W bits/s, and the total time duration is T s.

Now we invoke the crucial fact that space is a valuable resource in wireless networks. Since each transmission consumes a portion of the limited domain,

$$\sum_{b=1}^B \sum_{h=1}^{h(b)} \frac{\pi \Delta^2 r^2(b, h)}{16} \leq AWT.$$

The reason is that each packet of $W\tau$ bits with range r occupies $\pi \Delta^2 (r^2/16)$ m² for the duration of one slot of τ s, there are a total of T/τ slots, and the total area of the domain is A m².

We write the above as

$$\frac{1}{H} \sum_{b=1}^B \sum_{h=1}^{h(b)} r^2(b, h) \leq \frac{16 AWT}{\pi \Delta^2 H}.$$

Since r^2 is a convex function, we have

$$\left(\frac{1}{H} \sum_{b=1}^B \sum_{h=1}^{h(b)} r(b, h) \right)^2 \leq \frac{1}{H} \sum_{b=1}^B \sum_{h=1}^{h(b)} r^2(b, h).$$

Hence

$$\frac{1}{H} \sum_{b=1}^B \sum_{h=1}^{h(b)} r(b, h) \leq \sqrt{\frac{16 AWT}{\pi \Delta^2 H}}.$$

From (6) we obtain

$$\frac{BL}{T} \leq \sqrt{\frac{16 AWH}{\pi \Delta^2 T}}.$$

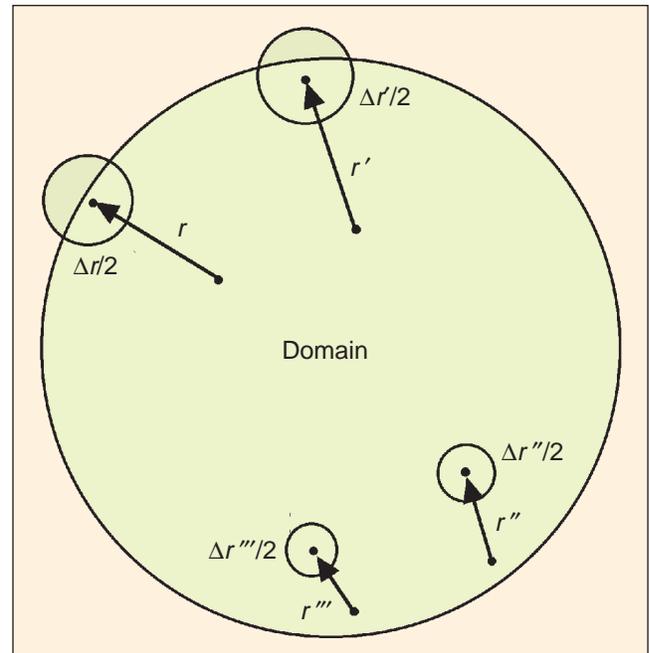


Figure 14. The exclusion zone around each receiver.

The result (5) then follows from (7).

This upper bound on the capacity, of order $O(W\sqrt{An})$, can in fact be achieved by arranging points in a regular grid (see [12]).

Hence $\Theta(W\sqrt{An})$ bit-m/s is a sharp characterization of the transport capacity. However, achieving it requires not only optimal *operation* of the network (i.e., optimal choice of transmission times, ranges, routes), but also optimal *location* of nodes, choice of traffic patterns, and so on. Note also that it is an upper bound even assuming the existence of a centralized controller that eliminates all conflicts and the resulting wasted slots.

The optimization of node locations and traffic patterns cannot be realized, of course. To address more common situations, a scenario can be considered where the nodes are randomly located in a disk and each node randomly chooses a destination. Assuming that all nodes employ the same range, it can be shown that the throughput which can be provided to every node is $\Theta(W/\sqrt{n\log n})$ bits/s [12]. Since the factor $1/\sqrt{\log n}$ is not terribly large, it follows that random networks may be nearly best.

One can also envisage other models for successful reception. Essentially, any such model's validity will depend on how receivers are constructed. Suppose that we construct receivers such that a packet is received successfully only when the signal-to-interference ratio exceeds some threshold. Assume also that the path loss of signals follows a $1/r^\alpha$ law for $\alpha > 2$. Then it is shown in [12] that the transport capacity is lower bounded by $c\sqrt{n}$ and upper bounded by $cn^{(\alpha-1)/\alpha}$. For α close to two, the case of inverse square-law loss, the bounds are nearly sharp.

One can imagine a much more sophisticated receiver design that ensures that the entire system operates at or near Shannon capacity. This design would require studying not only the capacity but also the complexity of decoding and encoding. This remains an area for future research.

Concluding Remarks

In the coming decades, the field of systems is headed for rapid change and greater technological vistas, à la the phase transition that occurred around 1960. Entire new technologies are within our purview. With the foundations laid over the past four decades, the opportunities are ours to grasp. As an example, we have illustrated the intrinsic problems in designing wireless networks.

Acknowledgement

This material is based on work partially supported by the Office of Naval Research under Contract N00014-99-1-0696, the Air Force Office of Scientific Research under Contract AF-DC-5-36128, by EPRI and the U.S. Army Research Office under subcontract to Cornell University, Contracts WO8333-04 and Army OSP 35352-6086, and by the U.S. Army Research Office under Contract DAAD19-00-1-0466.

References

- [1] S. Keshav, *An Engineering Approach to Computer Networking*. Reading, MA: Addison-Wesley, 1997.
- [2] T. Cover and J. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [3] C.C. Wu and D.P. Bertsekas, "Distributed power control algorithms for wireless network," in *Proc. 38th IEEE Conf. Decision and Control*, Phoenix, AZ, 1999, pp. 3556-3561.
- [4] S. Papavassiliou and L. Tassiulas, "Improving the capacity in wireless networks through integrated channel base station and power assignment," *IEEE Trans. Veh. Technol.*, vol. 47, pp. 417-427, 1998.
- [5] J. Rulnick and N. Bambos, "Mobile power management for wireless communication networks," *Wireless Networks*, vol. 3, pp. 3-14, 1997.
- [6] R. Yates and C. Huang, "Integrated power control and base station assignment," *IEEE Trans. Veh. Technol.*, vol. 44, pp. 638-644, 1995.
- [7] V. Jacobson, "Congestion avoidance and control," *ACM Computer Communi. Rev.*, pp. 314-329, August 1988. [Online] Available: <ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z>
- [8] V. Jacobson and M.J. Karels, "Congestion avoidance and control," in *Proc. SIGCOMM'88*, Aug. 1988.
- [9] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Networking*, vol. 1, pp. 397-413, Aug. 1993.
- [10] W.R. Stevens, *TCP/IP Illustrated*. Reading, MA: Addison-Wesley, 1996.
- [11] P. Gupta and P.R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, W.M. McEneaney, G. Yin, and Q. Zhang, Eds. Boston, MA: Birkhauser, 1998.
- [12] P. Gupta and P.R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inform. Theory*, vol. 46, pp. 388-404, Mar. 2000.
- [13] P. Gupta and P.R. Kumar, "Internets in the sky: The capacity of three dimensional wireless networks," *Commun. Inform. Syst.*, to appear.
- [14] F.A. Tobagi and L. Kleinrock, "Packet switching in radio channels: Part II-The hidden terminal problem in carrier sense multiple-access models and the busy-tone solution," *IEEE Trans. Commun.*, vol. COM-23, pp. 1417-1433, Dec. 1975.
- [15] N. Abramson, "The ALOHA system-another alternative for computer communications," in *Proc. Fall Joint Computer Conf. AFIPS Con.*, vol. 37, 1970.
- [16] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Standard 802.11-1997, 1997.
- [17] P. Karn, "MACA: A new channel access method for packet radio," in *Proc. 9th Computer Networking Conf.*, Sept. 1990, pp. 134-140.
- [18] V. Bharghavan, A. Demers, S. Shenkar, and L. Zhang, "MACAW: A media access protocol for wireless LANs," in *Proc. SIGCOMM'94 Conf. Communications Architectures, Protocols and Applications*, Aug. 1994, pp. 212-225.
- [19] S. Ramanathan and M. Steenstrup, "A survey of routing techniques for mobile communication networks," *Mobile Networks Applicat.*, vol. 1, no. 2, pp. 89-104, 1996.
- [20] D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1992.
- [21] D.P. Bertsekas and J.N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [22] J. Jubin and J.D. Tornow, "The DARPA packet radio network protocols," *Proc. IEEE*, vol. 75, pp. 21-32, Jan. 1987.
- [23] C.E. Perkins and P. Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," *Comput. Commun. Rev.*, vol. 24, no. 4, pp. 234-244, 1994.

[24] D.B. Johnson and D.A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*, T. Imielinski and H. Korth, Eds. Norwell, MA: Kluwer, 1996.

[25] C. Perkins and E. Royer, "Ad-hoc on-demand distance vector routing," in *Proc. IEEE Workshop Mobile Computing Systems and Applications*, New Orleans, LA, Feb. 1999.

[26] V.D. Park and M.S. Corson, "A highly adaptive distributed routing algorithm for mobile wireless networks," in *Proc. INFOCOM'97*, Kobe, Japan, April 7-11, 1997.

[27] P. Gupta and P.R. Kumar, "A system and traffic dependent adaptive routing algorithm for ad hoc networks," in *Proc. 36th IEEE Conf. Decision and Control*, San Diego, CA, 1997, pp. 2375-2380.

[28] J.G. Wardrop, "Some theoretical aspects of road traffic research," *Proc. Inst. Civ. Eng., Part 2*, pp. 325-378, 1952.

P.R. Kumar is Franklin Woeltge Professor of Electrical and Computer Engineering and a Research Professor in the Coordinated Science Laboratory at the University of Illinois, Urbana-Champaign. He obtained his B.Tech. degree from IIT, Madras, and the D.Sc. degree from Washington University, St. Louis. He was the recipient of the Donald P. Eckman Award of the American Automatic Control Council. He has presented plenary lectures at the IEEE Conference on Decision and Control, the SIAM Conference on Optimization, the SIAM Annual Meeting, and the Brazilian Automatic Control Congress. His current research interests are in wireless networks, wafer fabrication plants, machine learning, and financial economics.

Who wrote the following?

"..The emotional situation is much like that of an impecunious young man who has impetuously invited the lady of his heart to see a play, unmindful, for the moment, of the limitations of the \$2.65 in his pocket. The rapturous comments of the girl on the way to the theater would be very pleasant if they were not shadowed by his private speculation about the cost of the tickets."

For the answer, see page 118.

