1

Ad Hoc Networking An Introduction

Charles E. Perkins
Nokia Research Center

In recent years, mobile computing has enjoyed a tremendous rise in popularity. The continued miniaturization of mobile computing devices and the extraordinary rise of processing power available in mobile laptop computers combine to put more and better computer-based applications into the hands of a growing segment of the population. At the same time, the markets for wireless telephones and communication devices are experiencing rapid growth. Projections have been made that, by the year 2002, there will be more than a billion wireless communication devices in use, and more than 200 million wireless telephone handsets will be purchased annually. The rise of wireless telephony will change what it means to be "in touch"; already many people use their office telephone for taking messages while they are away and rely on their mobile telephone for more important or timely messages. Indeed, mobile phones are used for tasks as simple and as convenient as finding one's associates in a crowded shopping mall or at a conference. A similar transformation awaits mobile computer users, and we can expect new applications to be built for equally mundane but immediately convenient uses.

Much of the context for the transformation has to do with keeping in touch with the Internet. We expect to have "the network" at our disposal for the innumerable little conveniences that we have begun to integrate into our professional lives. We might wish to download a roadmap on the spur of the moment so that we can see what is available in the local area. We might wish to have driving suggestions sent to us, based on information from the global positioning system (GPS) in our car, using the services offered by various web sites. The combination of sufficiently fast and inexpensive wireless communication links and cheap mobile computing devices makes this a reality for many people today. In the future, the average traveler is likely to take such services for granted.

Today we see a great expansion in the production of technology to support mobile computing. Not only are the computers themselves getting more and more capable, but many new applications are being developed and wireless data communications products are becoming available that are much improved over those available in the past. The bandwidth now available to laptop computers over radio and infrared links is easily 10 to 100 times more than that available just ten years ago.

Such rapid technological advance has spurred equally impressive growth in mobile connectivity to the Internet. In the wired Ethernet domain, we have plug-and-play hardware and software so that laptop computers can be reconnected with ease according to the form factors of the local network outlets. The Internet is available around the world to those willing to make a dial-up connection to a local phone number. People are getting used to the advantages of having frequent and convenient Internet access. As a result, more and more network functionality will be taken for granted by typical laptop users.

As wireless network nodes proliferate and as applications using the Internet become familiar to a wider class of customers, those customers will expect to use networking applications even in situations where the Internet itself is not available. For instance, people using laptop computers at a conference in a hotel might wish to communicate in a variety of ways, without the mediation of routing across the global Internet. Yet today such obvious communications requirements cannot be easily met using Internet protocols. Providing solutions to meet such requirements is the subject of this book. The proposals to be described allow mobile computer users with (compatible) wireless communication devices to set up a possibly short-lived network just for the communication needs of the moment—in other words, an *ad hoc* network.

At the same time, there is a huge potential market for embedded network devices in our vehicles, our mobile telephones, and perhaps even in our toys and personal appliances. Surely the day is not far off when a typical child's doll will have a microprocessor and a remote control device and will depend on network access to interact with the home's television and computer games. Embedded networking could represent the "killer app" for wireless networks.

Anyone reading this book will agree that the modern age of networking represents one of the great achievements of humanity. We already take many aspects of it for granted. In particular, we often take for granted the infrastructure currently needed to support our vast networking enterprise. The things we do with our networks do not inherently depend on the network infrastructure; rather, having the infrastructure extends the reach of network applications immeasurably.

Once we have grown accustomed to the power of network communications and to accomplishing our daily tasks with the aid of applications that rely on networking, we will want the applications to be available at all times. In fact, many network researchers predict that some day in the not too distant future we will put our applications to use "anytime, anywhere," perhaps by way of the rapidly expanding satellite communications systems now under construction. The communications satellites girding the earth will complement the cellular (wireless) telephone infrastructure, which is itself growing even more rapidly in most developed countries.

Indeed, the authors of this book suggest that mobile computers and applications will become indispensable even at times when and at places where the necessary infrastructure is not available. Wireless computing devices should physically be able to communicate with each other, even when no routers or base stations or Internet service providers (ISPs) can be found. In the absence of infrastructure, what is needed is that the wireless devices themselves take on the missing functions.

In this introductory chapter, we consider some general topics that provide context for the rest of the chapters in this book. In the next section, we describe a general model of operation for ad hoc networks and some of the factors affecting the design decisions that various approaches have taken. In Section 1.2, we list a few of the commercial opportunities that may await vendors of wireless products when the necessary protocols are available. This will naturally include a look at some of the applications enabled by ad hoc networking. Following that, Section 1.3 will discuss some of the technical drivers for the resurgence of interest in ad hoc networking. The needs of military communications have been very influential in creating this renewed interest. Discussion of military ad hoc networking, however, is not included in that section because it is covered much more completely in Chapter 2. Because many of the approaches to ad hoc networks use variations on existing routing protocols, some very general comments about routing protocols are presented in Section 1.4. Finally, a capsule summary of each chapter in the book is presented in Section 1.5.

1.1 MODEL OF OPERATION

This book is concerned with ways (past and present) that wireless mobile computing devices can perform critical network topology functions that are normally the job of routers within the Internet infrastructure. Keeping track of the connections between computers is something so basic that a computer network, almost by definition, cannot exist without it.

There are many kinds of protocols available today that are supported by network infrastructure, either in a particular enterprise or in the Internet at large. These other protocols deserve consideration, but need adaptation before they can be useful within a network no longer connected to the Internet infrastructure. Some of them may not be appropriate for use when the infrastructure is unavailable; credit card validation and network management protocols come to mind.

As a matter of definition, an ad hoc network is one that comes together as needed, not necessarily with any assistance from the existing Internet infrastructure. For instance, one could turn on 15 laptop computers, each with the same kind of infrared data communications adapter, and hope that they could form a network among themselves. In fact, such a feature would be useful even if the laptops were stationary.

There are a bewildering variety of dimensions to the design space of ad hoc networks. We take a particular slice of that design space that should serve a large number of user requirements and yet allow discussion of a number of interesting and illuminating techniques. Besides *ad hoc networking,* similar techniques have been proposed under the names *instant infrastructure* [Bagrodia+ 1996] and *mobile-mesh networking* [SDT 1995].

Consider, for example, whether the range of wireless transmission should be large or small compared to the geographic distribution of the mobile wireless nodes. If all of the wireless nodes are within range of each other, no routing is needed, and the ad hoc network is, by definition, fully connected. While this might be a fortunate situation in practice, it is not a very interesting routing problem to solve. Plus, the power needed to obtain complete connectivity may be impractical, wasteful of battery power, too vulnerable to detection, or even illegal.

Thus, we discuss only proposals that offer solutions to the case in which some of the wireless nodes are not within range of each other. Combined with the lack of infrastructure routers, the restricted range of wireless transmission indicates the need for *multihop* routing.

As another example, we might suppose that wireless computer users could measure their relative positions and subsequently configure their laptop computers using the measured distances, so that the appropriate link information could be available at each mobile node. This would work, but it would not be very convenient. Worse yet, the link information would be likely to change whenever the users moved relative to each other. We are not interested in simplifying the problem space at the expense of user convenience, however, so we restrict our attention only to those proposals that provide automatic topology establishment (eschewing user configuration steps) and dynamic topology maintenance (enabling user mobility). In fact, we make the slight additional restriction of considering only proposals that are self-starting, except possibly for an enabling or mode setting step

performed by the user, who should be able to exert necessary controls over the performance of the ad hoc networking operation.

In this book, most of the discussion focuses on the interesting cases that have the following characteristics:

- The nodes are using IP, the Internet Protocol [Postel 1981a], and they have IP addresses that are assigned by some usually unspecified means.
- The nodes are far enough apart so that not all of them are within range of each other.
- The nodes may be mobile so that two nodes within range at one point in time may be out of range moments later.
- The nodes are able to assist each other in the process of delivering packets of data.

The discussion in this book focuses on the protocol engineering that underlies the establishment of the paths by which the ad hoc network nodes can communicate with each other. Thus, address autoconfiguration in particular, a very interesting subject, is largely absent from this book, but is ripe for exploration very soon.

As an example of a small ad hoc network, consider Figure 1.1 (taken from Chapter 3), illustrating a collection of eight nodes along with the links between them. The nodes are able to move relative to each other; as that happens, the links between them are broken and other links may be established. In the picture, MH_1 moves away from MH_2 and establishes new links with MH_7 and MH_8. Most algorithms also allow for the appearance of new mobile nodes and the disappearance of previously available nodes.

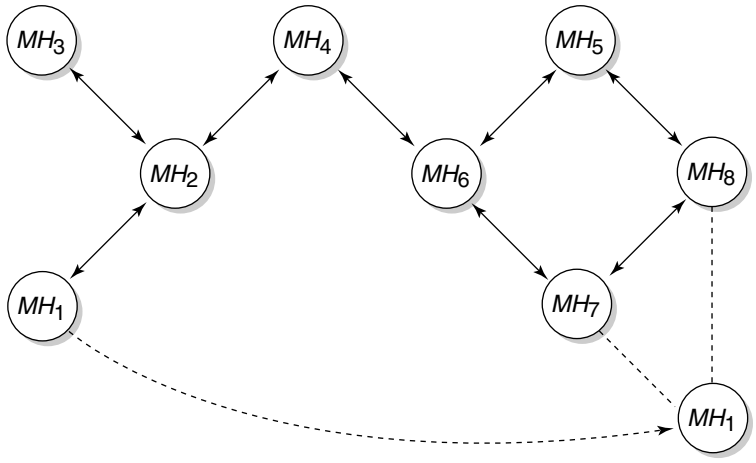

Figure 1.1. An Ad Hoc Network of Mobile Nodes

1.1.1 Symmetric Links

Some of the models considered in this book depend on the existence of symmetric communication links between the nodes in the ad hoc work. Unfortunately, wireless links in the real world do not necessarily conform to this assumption. This assumption of symmetry is made because routing in networks with unidirectional links is known to be quite difficult. Recent analysis has yielded a mathematical result of interest that characterizes this difficulty [Prakash 1999]. As it turns out, there are such networks in which two system-wide (i.e., all-node) broadcasts are needed for a source to find a route to a particular destination. On the other hand, not all networks with unidirectional links exhibit this characteristic. If the network has a sufficiently high degree of connectivity and relatively few unidirectional links, alternative routes comprising only symmetric links can usually be found.

There is another factor that mitigates the decision to ignore possible asymmetric routes. A unidirectional link is sometimes on the verge of failure anyway. In such cases, extending the basic ad hoc network protocol to deal with unidirectional links may cause less robust routes to be discovered, leading to early failure and the subsequent need for a new (and more complicated) route discovery cycle.

1.1.2 Layer-2 Ad Hoc Solutions

The ad hoc networking protocols in this book are mainly targeted at layer-3 operation. It is possible, in practically every case, to retool the protocol for use at layer 2. Doing so requires first that IP address fields be enlarged to contain 48 (or more) bits instead of 32, as needed for IP, because the IEEE MAC address is typically 48 or 64 bits long. This is not a problem, especially in view of the fact that such retooling will be needed anyway to enable the protocols to handle ad hoc networks of IPv6 addressable computers.

However, given the universal deployment of IP applications for networking today, every application eventually causes the communication subsystem to resolve an IP address into a neighboring layer-2 address, or else into the layer-2 address of a node in the neighborhood that can forward the application protocol data units (packets) toward the IP address of the desired application endpoint. When the route table at the application's source node has IP addresses of desired endpoints, IP forwarding naturally takes care of framing the data with a layer-2 header containing the layer-2 destination address of the next hop along a path toward the destination.

This happy circumstance evaporates, however, when the routing is based on layer-2 addresses. In that case, as a matter of consistent terminology, the route table is presumably indexed by layer-2 destination addresses.

That implies that the IP address of the destination has to be resolved to the layer-2 address of the destination, even for destinations that are multiple hops away. Then, unless the layer-2 route discovery is equipped with suitable layer-3 address information in appropriate extensions, additional broadcast discovery operations will be needed. If the layer-3 information is included for better performance, the entire operation can be viewed as a layer-3 route discovery anyway, albeit with an odd data structure for the route storage.

1.1.3 Proactive versus Reactive Protocols

One of the most interesting aspects of recent investigations concerns whether or not nodes in an ad hoc network should keep track of routes to all possible destinations, or instead keep track of only those destinations of immediate interest. A node in an ad hoc network does not need a route to a destination until that destination is to be the recipient of packets sent by the node, either as the actual source of the packet or as an intermediate node along a path from the source to the destination.

Protocols that keep track of routes for all destinations in the ad hoc network have the advantage that communications with arbitrary destinations experience minimal initial delay from the point of view of the application. When the application starts, a route can be immediately selected from the route table. Such protocols are called *proactive* because they store route information even before it is needed. They are also called *table driven* [Royer+ 1999] because we can imagine that routes are available as part of a well-maintained table.

However, proactive protocols suffer the disadvantage of additional control traffic that is needed to continually update stale route entries. The ad hoc network we are trying to support is presumed to contain numerous mobile nodes. Therefore, routes are likely to be broken frequently, as two mobile nodes that had established a link between them will no longer be able to support that link and thus no longer be able to support any routes that had depended on that link. If the broken route has to be repaired, even though no applications are using it, the repair effort can be considered wasted. This wasted effort can cause scarce bandwidth resources to be wasted and can cause further congestion at intermediate network points as the control packets occupy valuable queue space. Since control packets are often put at the head of the queue, the likely result will be data loss at congested network points. Data loss often translates to retransmission, delays, and further congestion.

As a result, *on-demand,* or *reactive,* protocols have been designed so that routing information is acquired only when it is actually needed. Reactive routing protocols may often use far less bandwidth for maintaining

the route tables at each node, but the latency for many applications will drastically increase. Most applications are likely to suffer a long delay when they start because a route to the destination will have to be acquired before the communications can begin.

One reasonable middle point between proactive and reactive protocols might be to keep track of multiple routes between a source and a destination node. This "multipath routing" might involve some way to purge stale routes even if they are not in active use. Otherwise, when a known broken route is discarded, one of the other members of the set of routes may be attempted. If the other routes are likely to be stale, the application may experience a long delay as each stale route is tried and discarded.

1.1.4 Multicast

Ad hoc networks are interesting to a large extent because of the challenge of maintaining a communication path between a source and a destination, even when some of the intermediate forwarding nodes are unable to continue participating in packet forwarding and must be replaced by nodes along another path. It turns out that maintaining paths between a single source and multiple destinations is somewhat more difficult, but not excessively so. Given the growing importance of multicast as a means to reduce the bandwidth utilization for mass distribution of data, and the pressing need to conserve scarce bandwidth over wireless media, it is natural that multicast routing should receive some attention for ad hoc networks.

It is an open question whether or not the multicast routing algorithm should be integrated with the routing algorithm used to establish communication paths between single endpoints. On the one hand, the problems may be sufficiently different so that trying to make a single routing algorithm serve for both may be unnaturally difficult. On the other hand, the problem of reestablishing paths caused by movement of intermediate points in a routing path or tree may dominate both situations. In this latter case, careful attention to path reestablishment for both contexts at once may be much easier than re-examining solutions in multiple disparate contexts, as would be necessary if multicast routing were unrelated to unicast routing.

1.2 COMMERCIAL APPLICATIONS OF AD HOC NETWORKING

In this section, we look at some of the potential applications for ad hoc networks that might provide the basis for commercially successful products. In fact, any commercially successful network application can be considered a candidate for useful deployment with nodes that can form ad hoc networks.

For example, users of nodes in an ad hoc network are likely to wish to transfer electronic mail. If some nodes in an ad hoc network offer web service, the other nodes that wish to make use of the service will still need to connect to the web server and support the usual HTTP traffic.

1.2.1 Conferencing

Perhaps the prototypical application requiring the establishment of an ad hoc network is mobile conferencing. When mobile computer users gather outside their normal office environment, the business network infrastructure is often missing. But the need for collaborative computing might be even more important here than in the everyday office environment. Indeed, the whole point of the meeting might be to make some further progress on a particular collaborative project. Given that today's project environments are heavily computerized, for projects in a very broad range of industries the need for being able to create an ad hoc network seems clear.

As it turns out, the establishment of an ad hoc network for collaborative mobile computer users is needed even when there may be Internet infrastructure support available. This results from the likely overhead required when utilizing infrastructure links, which might entail drastically suboptimal routing back and forth between widely separated office environments. Current solutions for mobile networking (e.g., Mobile IP [Perkins 1996]) are not well suited for efficiently supporting ad hoc networks, although the techniques are not wholly incompatible.

This interplay between the immediacy of ad hoc networks and the possible performance loss inherent in relying on Internet infrastructure routing is a recurring theme that should be considered in connection with each approach described in the chapters of this book.

1.2.2 Home Networking

As another example, consider the scenario that will likely result if wireless computers become popular at home. These computers will probably be taken to and from the office work environment and on business trips. It is quite possible that such computers will not have topologically related IP addresses, especially if they are connected at the offices of each parent or at the children's school. Keeping in mind the convenience that an unchanging IP address affords to the user, it would be nice to allow the various mobile computers to operate an ad hoc network in the home, even if the home maintains its own subnet with more or less permanently situated network nodes.

Add to this the fact that assigning multiple IP addresses to each wireless node for identification purposes would add an administrative burden (a job that most people do not want), and the alternative of deploying an

ad hoc network (automatically created as needed) seems more attractive. Ad hoc networking offers the prospect of reachability to all the nodes at home regardless of their "normal" point of attachment, which would otherwise be indicated by the network prefix that is part of every IP address. Furthermore, by using protocols such as Mobile IP, the nodes in the home ad hoc network can operate as if they were still connected to their standard computing environment *in addition* to participating (with higher performance) in the residential ad hoc network.

1.2.3 Emergency Services

When created at home or away from home at a meeting, an ad hoc network makes up for the lack of an existing Internet infrastructure. But, what about cases in which the existing infrastructure is damaged or out of service for other reasons? We are all familiar with situations in which loss of local power causes loss of electricity, and each year natural disasters wreak havoc with people's lives around the world. As the Internet grows in importance, the loss of network connectivity during such natural disasters will become an ever more noticeable consequence of the calamity. Furthermore, network applications will become increasingly important for emergency services, and thus it will be important to find ways to enable the operations of networks even when infrastructure elements have been disabled as part of the effects of a disaster.

Ad hoc networks can help to overcome network impairment during disaster emergencies. Mobile units will probably carry networking equipment in support of routine operations for the times when the Internet is available and the infrastructure has not been impaired. With the techniques and protocols in this book, emergency mobile units can greatly extend the usefulness of their networking equipment during times of lost infrastructure support. For instance, police squad cars and firefighting equipment can remain in touch longer and provide information more quickly if they can cooperate to form an ad hoc network in places not otherwise offering connectivity to the global Internet.

1.2.4 Personal Area Networks and Bluetooth

The idea of a personal area network (PAN) is to create a very localized network populated by some network nodes that are closely associated with a single person. These nodes may be attached to the person's belt or carried in a purse. More exotic visions of the future include virtual reality devices attached around the head and other devices more oriented toward the sense of touch. These devices may or may not need to have an attachment to the wide area Internet, but they will almost certainly need to communicate

with each other while they are associated with their users' activities. In this scenario, mobility is not the overriding consideration.

However, mobility suddenly becomes much more important when interactions between several PANs are needed. In other words, when people meet in real life, their PANs are likely to become aware of each other also. Since people usually do not stay in a fixed location with respect to each other for very long, the dynamic nature of this inter-PAN communication should be obvious. Methods for establishing communications between nodes on separate PANs could benefit from the technologies of ad hoc networks.

No current discussion about PANs would be complete without at least some mention of Bluetooth [Haartsen 1998]. Bluetooth is an emerging short-range radio technology targeted at eliminating wires between personal digital assistants (PDAs). If each PDA is equipped with a Bluetooth radio, it is possible for up to eight devices to organize themselves into what is called a *piconet,* with slotted communications controlled by a master. Bluetooth protocols at the physical and MAC layers are focused on saving battery power, as PDAs are much more useful if people do not have to be continually changing dead batteries.

When there are more than eight devices, Bluetooth requires that multiple piconets be formed. These piconets can be connected together into a *scatternet* if one of the slaves agrees to relay data between two of the masters. This suggests that the slave should have separate time slots in each piconet to reduce latencies for data transfers. It may even be necessary to form scatternets for the PAN that is associated with a single person. It is almost guaranteed that scatternets will be required for the interaction of multiple PANs.

1.2.5 Embedded Computing Applications

The world is full of machines that move, and future intelligent mobile machines will be able to process a great deal more information about the environment in which they operate. The "environment" itself will increasingly be a virtual one created by fixed and mobile computers. Some researchers [Weiser 1993] predict a world of *ubiquitous computing,* in which computers will be all around us, constantly performing mundane tasks to make our lives a little easier. These ubiquitous computers will often react to the changing environment in which they are situated and will themselves cause changes to the environment in ways that are, we hope, predictable and planned.

Many of these intelligent machines will be both mobile and connected by wireless data communications devices. The Bluetooth short-range radio device is expected to cost less than $5 within four years and to be incorporated into millions of wireless communications devices. Already, many

computers and PDAs are equipped with inexpensive wireless ports. These are being used for synchronizing data between machines owned by the same person, for exchanging virtual business cards, for printing out small files on local printers, and so on.

Ubiquitous intelligent internetworking devices that detect their environment, interact with each other, and respond to changing environmental conditions will create a future that is as challenging to imagine as a science fiction scenario. Our world will change so much that it is hard to predict the kinds of applications that might predominate. Security considerations must be taken into account, of course, to prevent unwarranted intrusions into our privacy and to protect against the possibility of impersonation by other people. However, these security matters are well known in other contexts, and they are not the particular subject of this book.

These capabilities can be provided with or without the use of ad hoc networks, but ad hoc networking is likely to be more flexible and convenient than, say, continual allocation and reallocation of endpoint IP addresses whenever a new wireless communication link is established. Furthermore, once we become used to having these simple and easily imagined features, it seems a sure bet that new applications will be invented. In fact, we may become so enamored of ad hoc computing that we will begin to presume the presence of appropriate support environments. These environments would be made available to ad hoc mobile nodes and could be considered a new form of micro-infrastructure. We might normally expect our ad hoc computers to have access to local information about temperature, light-switch controls, traffic information, or the way toward a water fountain [Hodes+ 1997]. Infrastructure elements that make their information available by way of standard TCP/IP client–server applications could operate in dual mode so that they participate in ad hoc networks as well, depending on the circumstances.

1.2.6 Sensor Dust

Recent attention has been focused on ideas involving the possibility of coordinating the activities and reports of a large collection of tiny sensor devices [Estrin+ 1999, Kahn+ 1999]. Such devices, cheap to manufacture and able to be strewn about in large numbers of identical units, could offer detailed information about terrain or environmental dangerous conditions. They might be equipped with positional indicators; alternatively, positional information could be inferred for network information such as the number of hops between various sensors and a well-known data collection node.

Such sensor network nodes have two characteristics that will strongly influence the design of networks to facilitate data acquisition:

- Once situated, the sensors remain stationary.
- Population may be largely homogeneous.
- Power is likely to be a scarce resource, so sophisticated communications scheduling will be important (see Section 1.3.2 and Chapter 10).
- In fact, the lifetime of the battery may define the node's lifetime.

As an example, suppose some hazardous chemicals were dispersed in an unknown manner because of an explosion or some other sort of accident. Instead of sending in emergency personnel who might be subjected to lethal gas and forced to work in unwieldy protective clothing, it would be better to distribute sensors containing wireless transceivers (perhaps by dropping them from a low-flying plane). The sensors could then form an ad hoc network and cooperate to gather the desired information about chemical concentrations and identification. Military applications are also of great interest.

1.2.7 Automotive/PC Interaction

In a somewhat different vein, consider the possible uses of ad hoc networks between automotive computers and laptops or PDAs that may accompany us as we travel in our cars. Suppose that when we start the car, we have an indication on our display that there may be a mechanical difficulty that needs attention. We may expect our mobile telephone to provide a link between our car and some local repair service advertised in a local service directory available by way of our PDA. Alternatively, once the service needs from the car's network have been loaded onto a laptop computer that we happen to be carrying, the laptop might take charge of finding directions to the repair shop. A laptop computer is likely to have a superior display monitor, and we might even want to take a little extra time to evaluate reports from various satisfied or dissatisfied customers before selecting one of several repair shops. The selection of display device and the urgency of making the decision are still matters for personal attention, but our computing devices can make matters much easier by cooperation and automatic discovery of relevant information for our consideration.

There are other interesting interactions that can be programmed between automobiles and their occupants. Positional information will likely be available to drivers and passengers, allowing those in the back seat to easily find answers to their typical questions. That might become a concern in communications between cars that are trying to arrange a meeting place. Wireless communication between cars could become the logical successor to citizen's band (CB) radio. Indeed, with both audio and video over wireless within reach of today's wireless technologies, countless possibilities for both useful and frivolous communications are easily imaginable. Browsing the web pages of nearby cars might become a new national pastime.

1.2.8 Other Envisioned Applications

Once they become conveniently available by way of widely deployed dynamic routing protocols as described in this book, ad hoc networks might be useful in many ways we cannot fully understand given our current experience. For instance, university campuses might well become large ad hoc networks as students and faculty learn to rely on their handheld and laptop computers for their communication and computing needs. Messaging and browsing could be managed either by available wireless infrastructure or by ad hoc connectivity, according to whatever is most convenient at the moment.

Similarly, at hospitals busy doctors and nurses will want to rely on the administrative infrastructure at times and to instantiate direct links outside the infrastructure at other times. Visiting staff and paramedics will need to confer with residents and transfer data to and from patient equipment. These operations will often be facilitated by infrastructure support, but the same communications needs can sometimes be met without interaction with the infrastructure.

Viewed from another perspective, however, searching for applications to justify the development of ad hoc networks may be putting the cart before the horse. What really matters is making communications technology useful for people everywhere regardless of the nature or availability of backbone infrastructure. We do not require a killer app for ad hoc networking any more than we needed a killer app for the success of the backbone Internet.

When people are within wireless range of each other, it becomes merely an unfortunate artifact of decades-old technology that their communication devices require remote infrastructures. With ad hoc networking, local communications can instead depend only on local communications channels and transmission technologies and protocols. As this localized technology gains significant mindshare, it is to be hoped that sufficient spectrum will be allocated for local use, given that the current ISM spectrum bands will not serve future needs. Furthermore, we can hope that such localized communications will rekindle the public's interest in reasserting its right to use the spectrum, which is supposed to be a resource managed for the public good. The rise of wireless cellular telephony, combined with ad hoc networks, could provide a new paradigm of public use of the public airwaves.

1.3 TECHNICAL AND MARKET FACTORS AFFECTING AD HOC NETWORKS

Nodes in an ad hoc network are often assumed to have IP addresses that are preassigned, or assigned in a way that is not directly related to their current position relative to the rest of the network topology. This differs

substantially from the way that IP addresses are assigned to nodes in the global Internet. Routing within today's Internet depends on the ability to *aggregate* reachability information to IP nodes. This aggregation is based on the assignment of IP addresses to nodes so that all the nodes on the same network link share the same *routing prefix*. In practice, good network administration requires that networks that are nearby should have similar prefixes. Classless Inter-Domain Routing (CIDR) [Rekhter+ 1993] has been very effective in helping to reduce the number of routing prefixes that have to be advertised across the Internet, thus enabling today's router hardware to continue to maintain the Internet's global addressability. Using CIDR, when the prefixes themselves can be aggregated, reachability to all the networks within a site can be described by advertising a prefix that is itself the initial part of all the (longer) prefixes for the networks within the site. The process of aggregation can be iterated if network administrators for nearby sites cooperate to use networks with prefixes that share a common initial bit string (i.e., a common smaller prefix).

This creates a hierarchy of network prefixes—smaller prefixes that fit at higher levels of the hierarchy. Reachability to all nodes within the hierarchy can be described by advertising a single, smallest routing prefix. This drastically reduces the amount of routing information that has to be advertised and provides the necessary economy for the Internet to continue to grow. We can say that aggregating routing information is the key to Internet *scalability*.

With ad hoc networks, however, such aggregation is typically not available. Some proposed methods attempt to reintroduce aggregation by controlling the IP addresses of the mobile nodes, but this requires that the IP addresses (and, subsequently, routing information relevant to the mobile node) be changed depending on the relative movement of the node. It is not at all clear that the benefit of improved aggregation is worth the cost of complicated re-addressing and route table revisions. Thus, the scalability afforded by aggregation within the Internet is not likely to be available for ad hoc networks. This means that there may be major limitations on the viability of ad hoc network algorithms for extremely large populations of mobile nodes. The limits will also depend on the relative speed of movement between the mobile nodes. More movement means more maintenance so that the available routing information remains useful. In the face of uncontrolled increases in node mobility, any ad hoc routing algorithm will eventually require so much route maintenance that no bandwidth remains for the transmission of data packets [Corson+ 1996].

1.3.1 Scalability

Because ad hoc networks do not typically allow the same kinds of aggregation techniques that are available to standard Internet routing protocols,

they are vulnerable to scalability problems. In particular, loss of aggregation leads to bigger route tables. There are ways to maintain aggregation for ad hoc networks, and some of them are discussed in Chapter 4. Aggregation can be very easily used for mobile networks, but the aggregation and addressing used for routing within the structure are not IP based. Consequently IP-based ad hoc protocols often must use additional memory for storing the route tables and processor cycles for searching them.

Node mobility introduces other kinds of scalability problems for ad hoc networking protocols. Since the routing changes as the nodes move, control messages have to be sent around the network to represent current connectivity information. These control messages are likely to be transmitted more often if the nodes move more quickly relative to each other, because then links will be lost or established more often between the nodes. This usually happens be true unless the movements of the mobile nodes are highly correlated.

The increased number of control messages places additional load on the available bandwidth, which is usually already a constraining factor for communication between wireless nodes. Thus, ad hoc protocols are typically designed to reduce the number of control messages, often by maintaining appropriate state information at designated mobile nodes. The downside of maintaining state information, however, is that it can become stale; the only cure for updating stale information is the introduction of more control messages.

Depending on the details of the algorithm, transmission of control messages may cause undesirable loads on the individual processing elements as well as on the available network bandwidth. For instance, protocols that cause a recomputation of the entire network topology whenever a new routing update is received may be subject to long convergence times whenever a node makes or breaks a link with one of its neighbors. The data in such a route update must be processed in far less time than the average time between network events caused by node mobility; otherwise, the ad hoc network may never stabilize. Route instability is a likely cause of routing loops, which can cause unnecessary bandwidth consumption. By Murphy's Law, such problems always arise at exactly the moment when they are least welcome or at exactly the time when communication is most critical.

Algorithms for ad hoc networking must be carefully evaluated and compared for their relative scalability in the face of node population growth and increased node mobility. If maximum values are known for those numbers, it is reasonable to calculate how many control messages may be required to manage the ad hoc network and to compare the total traffic due to control messages against the total bandwidth availability. As long as the control traffic takes up a manageable proportion of the overall bandwidth, a candidate protocol can be considered acceptable. Similarly, the time taken for convergence should be calculated for a given maximum value for node mobility if such a value is known.

1.3.2 Power Budget versus Latency

In many kinds of ad hoc networks, the mobile nodes operate on battery power. There are two ways that they do this. First, they might transmit data to a desired recipient. This use of battery power is not part of the overhead of ad hoc networking. Second, a mobile node might offer itself as an intermediate forwarding node for data going between two other nodes in the network. Providing such a service is likely to be costly in terms of power consumption, but without the availability of such forwarding nodes there can be no ad hoc network.

There are interesting questions about when a node should or should not forward traffic. For instance, perhaps a node with full battery power should be more willing to forward data for its neighbors than a node whose battery power is almost depleted. Nodes with reduced power might limit their activities to transmitting and receiving only emergency or high-priority messages. Server nodes may attempt to reserve bandwidth for sourcing data and rely on their other neighbors for enabling route establishment between other endpoints. Other nodes may attempt to "freeload" from their neighbors, taking advantage of their forwarding services without offering anything in return. If this potential bad behavior is of concern, steps should be taken to isolate the offending nodes and to offer forwarding services only to cooperative nodes. Detecting the offense, however, can be problematic or impossible. For instance, the behavior of any leaf node might be indistinguishable from that of a selfish node.

The behavior of the nodes in the ad hoc network given their power budget is likely to affect the ease with which routes can be established between endpoints. If it takes more control messages to find or maintain a route, communication latency might be increased. If route information is periodically transmitted throughout the network, further demands will be placed on the power budget of each individual node. However, the more route information that is made available, the more likely it is that a good route between endpoints can be found as soon as it is needed without any additional control operations. This is especially true as the frequency of periodic transmissions increases, because then it is more likely that existing (cached) route information is still valid. The tradeoff between frequency of route update dissemination and battery power utilization is one of the major engineering design decisions for ad hoc network protocols.

1.3.3 Protocol Deployment and Incompatible Standards

Wireless ad hoc network protocols may be susceptible to forces working against standardization. If the experience of the IEEE 802.11 committee is any guide, we can expect to see a plethora of engineering solutions, all incompatible and all solving a slightly different part of the problem. In the

IEEE, this was in part due to the huge design space for physical wireless channel access and coding techniques. Unfortunately, the design space for wireless routing protocols is also impressively large, as will become evident from the discussion in Section 1.4 and in the rest of this book. Furthermore, the need for standardization in routing protocols at the network layer is almost as great as the need at the lower protocol layers.

At least when neighbors share the same physical medium and method of utilizing it (e.g., MAC layer), there is hope for communication within a local neighborhood. Communication between nodes not in the same neighborhood (e.g., not sharing the same physical medium) will not be possible unless the nodes also agree on some higher-level protocol by which they can interchange connectivity information about links outside their respective neighborhoods. Unless a miracle happens (e.g., the IETF `manet` working group is able to promulgate a widely deployed ad hoc networking protocol), ad hoc networks will gain momentum only gradually because users will have to load software or take additional steps to ensure interoperability. But sooner or later, just as with the IEEE 802.11 standardization process, some useful standards are bound to emerge.

1.3.4 Wireless Data Rates

One of the biggest obstacles to the adoption of ad hoc networks may be reduced data rates—the same problem that slowed the adoption of wireless computing during the last decade. We can typically observe an order of magnitude difference in the speed of wired and wireless networks. For instance, while many enterprise users are accustomed to 100 Mbit/sec from the local Ethernet, wireless users must struggle to get a reliable 10 Mbit/sec over the air: 1 to 2 Mbit/sec is much more common.

The overall effect tends to be that wireless computers are no longer *general purpose*. The wireless user has to be careful not to invoke applications that require a lot of bandwidth. As many of today's applications involve transactions over the Web, it may not be so easy for the user to avoid this problem. At any moment, the next hyperlink selection may attempt to load some beautiful dancing alphabetic letters on fire, at great cost in bandwidth or at great cost in frustration as the user tries to figure out what went wrong. It is often hard to tell whether the network itself has failed or whether the implicitly selected image is huge and is tying up the available communication paths.

A related problem has to do with the higher error rates experienced by wireless media in comparison with today's wired network media. Because TCP was mostly designed to classify a lost packet as a sign of network congestion, it tends to respond poorly when data errors cause a packet to be lost or dropped. Thus, the hapless wireless user is stuck with an

even greater performance loss whenever transient noise or obstacles cause a temporary increase in errors. This problem, while receiving quite a bit of recent attention, is not close to any solution at all, much less a widely deployed one. Furthermore, there are indications that TCP itself performs even worse than expected across multiple wireless hops.

Several IETF efforts have been aimed at identifying a working group charter for reaching a solution, but as of this writing no working group has been formed. There is no agreement about where to start searching for a solution to consider for standardization. As existing Internet applications use TCP, and as the same applications are likely to find use in ad hoc networks, TCP will probably be the first transport protocol of interest for ad hoc networks. Thus, the problems with wireless error rates are of central importance for ad hoc network designers.

1.3.5 User Education and Acculturation

Many of the obstacles to widespread deployment of wireless data devices stem from user education and acculturation. Menu selection, avoidance of web pages with huge image files, and alphanumeric data entry on small user input devices are all problematic for the uninitiated. However, these obstacles are slowly disappearing as PDAs such as the wireless Palm Pilot, WAP terminals, and iMode devices continue to grow in popularity. User input restrictions are mitigated either by special input modalities or by simplified menu selection from specially engineered web pages. The storage of convenient user profiles on the Web also helps reduce the requirements for data entry on bandwidth- and space-constrained devices.

Sometimes the obstacles to user acceptance are the most prosaic and mundane features. Current wireless devices depend on an external or exposed antenna for reliable operation. The antenna turns out to be expected and yet a pet peeve and object of frequent irritation. It is so much expected that some wireless telephones have a false antenna that can be pulled out by the user in order to "improve" the perceived voice quality. However, any such exposed device is constantly in the way and vulnerable to breakage.

1.3.6 Additional Security Exposure

As with any wireless communications, traffic across an ad hoc network can be highly vulnerable to security threats. What distinguishes an ad hoc network is that the additional threats extend even to the basic structure of the network. Thus, existing techniques for securing network protocol transactions for wired networks (see Section 1.4) must also be applied to the techniques of ad hoc networking. Unfortunately, this is easier said than done.

In the first place, security for routing protocols almost always depends on proper distribution of some key that allows the creation of unforgeable credentials. Designing secure key distribution in an ad hoc network might be a frightening prospect. Any reliance on a certificate authority seems doomed from the outset, for the same reason that reliance on any centralized authority is problematic. Centralization is antithetical to ad hoc networking, if not outright contradictory.

Beyond that, however, there are additional problems with the increased packet sizes required by authentication extensions. It is likely that the more secure a protocol is made to be, the slower and more cumbersome it will become. This combination of poor performance and tedious, inconvenient key distribution and configuration probably means that ad hoc networks will remain characteristically insecure. Diffie-Hellman key exchange techniques will help establish some temporary security between particular endpoints, but they are vulnerable to *man-in-the-middle* attacks that are hard to defeat in an ad hoc environment.

1.3.7 Spotty Coverage

Gaps in wireless coverage are both a problem and an opportunity for peer-to-peer devices capable of forming ad hoc networks. On the one hand, the coverage gaps make it unlikely that people in the affected regions will invest in most existing wireless devices, because such devices depend on infrastructure support for their operation.

On the other hand, people who are used to wireless connectivity to the Internet will often have to travel through regions of poor connectivity. These are likely to be the people who are motivated to put ad hoc network products into place so that they can continue to use local intercommunications even when the wide area infrastructure is inoperative.

1.4 GENERAL COMMENTS ON ROUTING PROTOCOLS

Each node in an ad hoc network, if it volunteers to carry traffic, participates in the formation of the network topology. This is quite similar to the way that intermediate nodes within the Internet, or within a corporate intranet, cooperate to form a routing infrastructure. Routing protocols within the Internet provide the information necessary for each node to forward packets to the next hop along the way from the source to the destination.

This observation motivates attempts to adapt existing routing protocols for use in ad hoc networks. Routing protocols are self-starting, adapt to changing network conditions, and almost by definition offer multihop paths

across a network from a source to the destination. From this description, it is easy to see why we might wish to manage topology changes in an ad hoc network by requiring all forwarding nodes to operate some kind of routing protocol.

The wired Internet uses routing protocols based on network broadcast, such as OSPF [Moy 1997]. OSPF is an example of a *link-state* protocol, so named because each node gathers information about the state of the links that have been established between the other nodes in the network. Dijkstra's shortest path first (SPF) algorithm [Dijkstra 1959] can be used to construct routes as needed between sources and destinations for which the link-state information is available. Traditional link-state protocols may not be suitable for highly dynamic networks because of the relatively large bandwidth required to maintain a current view of the network state. Experiments have shown that OSPF can consume a very substantial percentage of the available bandwidth while trying to maintain adequate link information for routing in a typical military ad hoc network (Chapter 2, [Strater+ 1996]).

However, there are alternatives to table-driven link-state routing approaches. Some current link-state algorithms do not require all nodes to have identical link-state information and route selection algorithms, and generate routes on demand. Examples of these appear in Chapters 4 and 10.

Instead, nodes distribute link-state information (usually in the form of a statistical, not an instantaneous, characterization) only when there has been a significant change. Furthermore, this link-state information is distributed only to nodes that might be affected by the change—for example, those in close proximity to the place where the change occurred or, in a clustered network, those within the same cluster. It is not necessary for all nodes to receive all link-state information intended for them in order to compute routes.

Routing in multihop packet radio networks was based in the past on shortest-path routing algorithms [Leiner+ 1987], such as the Distributed Bellman-Ford (DBF) algorithm [Ford+ 1962]. DBF algorithms typically store very little information about links that are not directly connected to the node running the algorithm. For each destination, the node typically stores just a single route table entry, containing among other things the next hop toward that destination. DBF algorithms are also known as *distance-vector* (DV) algorithms because the route table entry for a destination contains a *metric,* which is often just the *distance* from the node to the destination, as well as the next hop (or *vector*) toward the destination.

Distance-vector algorithms have many advantages. They are easy to program. UC Berkeley offered a free implementation of the Routing Information Protocol (RIP) [Malkin 1994], which is easy to understand and modify. However, these algorithms can suffer from very slow convergence (the "counting to infinity" problem). This is because neighboring nodes can

confuse each other by passing stale information back and forth, incrementing the distance to a destination each time, until the recorded distance exceeds the maximum allowable (i.e., infinity). Techniques such as *split-horizon* and *poisoned-reverse* are used to mitigate the danger of counting to infinity; unfortunately, they do not eliminate it. Early implementations tried to solve the problem by defining infinity to be 15, but that clearly was a very short-term solution.

Compared to link-state algorithms, distance-vector algorithms use less memory and, since state information is not stored at every node, more localized updates. However, when the distance to a particular destination changes, the effect can still ripple through every routing node in the network.

DBF-like protocols incur large update message penalties, although not as large as those incurred by link-state protocols. Attempts to fix some of the shortcomings of DBF, such as Destination-Sequenced Distance-Vector routing (DSDV) (see Chapter 3) have been proposed. However, synchronization and extra processing overhead are common in these protocols. Other protocols that rely on the information from the predecessor of the shortest path solve the slow convergence problem of DBF (e.g., [Cheng+ 1989, Garcia-Luna-Aceves 1993]). The processing requirements of these protocols may be quite high because of the way they process the update messages.

Scaling by grouping nodes into clusters, abstracting state information for the clusters, and selecting routes according to this abstracted state form a very old concept in networking, starting in the early 1970s with MacQuillan, Kamoun, and Kleinrock. Moreover, this hierarchical structure of nested clusters yields a natural addressing structure, such that a node's address is based on the labels of its ancestral clusters (i.e., the clusters in which it is contained). Routing within such a hierarchical clustering structure is what is usually termed *hierarchical routing*. Hierarchical routing addresses the inefficiency of globally propagating local topology information. Most schemes using hierarchical routing are designed to avoid routing only through the clusterhead, which otherwise would introduce overhead not found in flat routing schemes. One approach to hierarchical routing is the Landmark Hierarchy [Tsuchiya 1988], which is discussed in Chapter 4 along with many other useful variants.

Other routing protocols have been adapted for use with ad hoc networks, notably source routing (see Chapter 5) and link-reversal algorithms (see Chapter 8). Ad hoc networks can be viewed as the "acid test" of network protocol design, and so they are likely to continue to be of major interest. It may turn out that routing protocols designed for ad hoc networks can be adapted to greatly improve the scalability of routing protocols designed for use in the global Internet. That would be an enormous payoff for ad hoc network research.

1.5 DESCRIPTION OF THE MATERIAL PRESENTED

The actual technical material follows this introductory chapter. Most of the chapters focus on a particular ad hoc networking protocol. It is fortunate that we can include in this book the descriptions presented in detail by the person or persons responsible for the creation of the protocol. Frequently, simulation results accompany the description of the algorithms and protocols. Each chapter and protocol has been selected to offer the reader the most diversity of opinion. The reader must decide separately how to make the comparisons.

Chapter 2—A DoD Perspective on Mobile Ad Hoc Networks

Chapter 2 describes the background of ad hoc networking from the military point of view and gives some ideas about how ad hoc network protocols are important for the armed forces. During combat, every design parameter for the communications network is stressed to, and sometimes beyond, the breaking point. There are many reasons for failure of military networks beyond the workaday problems of traffic congestion and crashing network nodes.

The military needs for communication engendered some of the earliest and best research into ad hoc networks, so this chapter seems almost mandatory for the reader to gain the fullest understanding of the subject.

Chapter 3—DSDV: Routing over a Multihop Wireless Network of Mobile Computers

My own involvement with ad hoc networks began with the design of the Destination-Sequenced Distance-Vector (DSDV) protocol. This protocol was designed in an ad hoc fashion; with Pravin Bhagwat, I started with Berkeley's routed and fixed problems one by one as they cropped up. I was surprised when Pravin could not find any previous work that included the idea of a destination sequence number to eliminate the counting to infinity problem. Because we thought we had found a very easy solution to a well-known problem, and because the solution was also easy to program, we decided to push it as far as we could. This was quite a while before we thought about doing anything "on-demand," so DSDV looks primitive by comparison with modern protocols.

Chapter 3, on DSDV, is presented with very few updates from the original.[1] I decided to include it in this book mainly for historical purposes. It was very instructive to build the protocol, and the protocol has

[1] In *Mobile Computing* edited by Tomasz Imielinski and Henry F. Korth. Norwood, MA: Kluwer, 1996—Chapter 6, pp. 183–206, by Perkins and Bhagwat.

some features that may be considered innovative even today. There have been numerous comparisons with DSDV, and I often think that is so because DSDV is so easy to beat in performance in many applications. Even now, papers propose new slants on how to outperform DSDV. Nevertheless, DSDV is still possibly a good contender for scenarios in which almost all nodes are mutually involved in communications with almost all other nodes and in which the mobility factor is neither very low nor very high.

Chapter 4—Cluster-Based Networks

Chapter 4 emphasizes routing within networks that have been organized according to cluster-based control structures, where the purpose of a cluster-based control structure may be for transmission management, backbone formation, or routing efficiency. The algorithms for forming and using these clusters and the structure and interconnectivity of the resulting clusters depend strongly on the purpose for which they are intended. The chapter presents various approaches, grouped according to the three purposes here and to show similarities and differences in the approaches and the advantages and disadvantages of each. In addition, a fine list of references is provided.

Chapter 5—DSR: The Dynamic Source Routing Protocol for Multihop Wireless Ad Hoc Networks

The Dynamic Source Routing (DSR) algorithm is another innovative approach to ad hoc networking, whereby nodes communicate along paths stored in source routes carried along with the data packets. DSR explores the many advantages of source routing and enjoys the benefits of some of the most extensive testing and deployment of any of the protocols in this book. It is one of the purest examples of an *on-demand* protocol, in which all actions are taken only when a route is actually needed. While reading Chapter 5, the reader should think about the ways that the additional path information in source routes can be applied to get a fuller description of the network topology even as it is changing over time.

Chapter 6—AODV: The Ad Hoc On-Demand Distance-Vector Protocol

Having experience with distance-vector routing and sequence numbering from DSDV, several researchers and I took up the task of shaping an *on-demand* version of that distance-vector protocol, called Ad Hoc On-Demand Distance-Vector (AODV). We were confident that nobody else had used that particular acronym, so we were safe at least on that count.

AODV offers a pure distance-vector approach to the problems of ad hoc networking, which means reduced memory and processing requirements. Because it acquires and maintains routes only on demand, the control traffic is reduced compared to most table-driven protocols. AODV takes care to cache routes, as long as the routes are likely to remain valid, to reduce unnecessary route acquisition; on the other hand, routes are purged soon after they no longer appear to be useful so that no additional control traffic will be wasted to maintain unused route information. AODV also offers multicast that benefits from the same route caching algorithms, address aggregation, some quality of service, and address autoconfiguration.

Over the years, AODV has benefited (as have many of the other protocols in this book) from the storehouse of knowledge built up in the IETF `manet` Working Group. As each Internet draft is published, new discussion ensues, identifying problems and possible solutions. This new and cooperative way of defining network protocols has been a very satisfying experiment in collaborative design, with the principal protocol designers borrowing the expertise of some of the best talents available.

Chapter 7—ZRP: A Hybrid Framework for Rerouting in Ad Hoc Networks

The Zone Routing Protocol (ZRP) takes a fresh yet time-tested approach to protocol improvement by constructing a way to hybridize table-driven protocols (such as DSDV) with on-demand protocols. ZRP uses zones that are similar to clusters, but instead of hierarchical routing between clusters being used, special border nodes are dynamically selected that connect adjacent zones. A zone radius parameter dynamically adjusts the size of the zone, in terms of the number of hops, as the network topology changes. A different routing protocol can be used between zones as compared to the one used within a zone. A proactive scheme is used inside the zone, and outside the zone routes are discovered only reactively.

This approach is almost guaranteed to find a happy medium between the two extremes that exhibits improved properties. The chapter first describes how the hybrid protocol can be parameterized to yield the extreme design points and then defines metrics for taking measurements. It is instructive to think about other extremes and how they might be hybridized to produce still other variants.

Chapter 8—Link-Reversal Routing

The Temporally Ordered Routing Algorithm (TORA) is the newest descendent of several *link-reversal* protocols derived from the original Gafni-Bertsekas algorithm, which is also described in this chapter. Creating a

route between nodes is accomplished by building a directed acyclic graph (DAG). A packet is routed on the basis of information at each node (router) in the network, and routes are selected through evaluation of a sophisticated *height* function.

This innovative design will surely convince the reader that there is no realistic limit to the creative protocol design that can be brought to bear on the problems of ad hoc networking. I have often wondered if there is a simple transformation that would reduce TORA to, for instance, AODV or DSDV. Alternatively, if the height function were augmented with a source route, then perhaps TORA could also be considered a generalization of DSR. It is a fascinating mental exercise to imagine such protocol transformations and the generalities that must be extracted to produce each such specialization. I also believe that TORA can be specialized to produce each preceding descendant from the original Gafni-Bertsekas algorithm.

Chapter 9—The Effects of Beaconing on the Battery Life of Ad Hoc Mobile Computers

Chapter 9, on Associative Bit Routing (ABR), focuses on battery life, but I think that ABR is an example of yet another diverse design point for ad hoc network protocols. Like each preceding approach, ABR is a natural development of a protocol based on a sensible and intuitive model for route selection. Because battery life is improving at a rate far slower than that of improvements in processor speed and memory density, we will eventually face the prospect that caring for battery life may become the prime justification for many protocol design features. This is especially true for applications with tiny and very inexpensive devices like sensor dust, as described in section 1.2.6.

Chapter 10—Bandwidth-Efficient Link-State Routing in Wireless Networks

The most well-known link-state algorithms, such as the original ARPANet and the later OSPF and IS-IS, are table driven and require complete link-state information. Derived from these algorithms is a new type of link-state algorithm—a partial link-state algorithm that is not table driven.

Chapter 10, the last technical chapter, presents STAR as an example of a partial link-state algorithm. It also explores several other design innovations that enable maintenance of multiple routes between source and destination. The protocol in this chapter also has been extended to manage additional link-state parameters such as bandwidth availability and queue length; the latter quantity is directly related to average delay. In keeping with our theme, STAR bears little resemblance to the other

protocols in this book. Thus, this last protocol chapter offers further convincing proof that the design space for ad hoc network protocols is huge indeed.

Chapter 11—Summary and Future Work

Finally, I close the book with some modest observations about possible futures for the field of ad hoc networking. I hope that the reader will be gentle in his or her criticism, for my attempt at identifying an ad hoc future is only one possibility out of the huge space of design and contingency that awaits all designers of ad hoc network protocols.

I hope that this book will delight and invigorate those who read it with the many divergent intuitions and creative energies that emerge from each chapter. At the end of the last chapter, I have made a list of relevant mailing lists and some other resources for participation. This, combined with the references in each chapter, should provide a resource with something to offer everyone.

References

[Bagrodia+ 1996] R. Bagrodia, M. Gerla, L. Kleinrock, J. Short, and T.-C. Tsai. *A Hierarchical Simulation Environment for Mobile Wireless Networks*. Technical report, Computer Science Department, University of California at Los Angeles, 1996.

[Cheng+ 1989] C. Cheng, R. Riley, S.P.R. Kumar, and J.J. Garcia-Luna-Aceves. A Loop-Free Extended Bellman-Ford Routing Protocol without Bouncing Effect. *ACM Computer Communication Review* 19(4):224–236, May 1989.

[Corson+ 1996] M.S. Corson, J. Macker, and S. Batsell. Architectural Considerations for Mobile Mesh Networking. In *Proceedings of the IEEE Military Communications Conference (MILCOM '96),* October 1996.

[Dijkstra 1959] E.W. Dijkstra. A Note on Two Problems in Connection with Graphs. *Numerische Math.* 1:269–271, 1959.

[Estrin+ 1999] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Scalable Coordination in Sensor Networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM),* August 1999, 263–270.

[Ford+ 1962] L.R. Ford, Jr., and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, N.J., 1962.

[Garcia-Luna-Aceves 1993] J.J. Garcia-Luna-Aceves. Loop-Free Routing Using Diffusing Computations. *IEEE/ACM Transactions on Networking* 1(1):130–141, February 1993.

[Haartsen 1998] J. Haartsen. Bluetooth—The Universal Radio Interface for Ad Hoc Wireless Connectivity. *Ericsson Review* (3), 1998.

[Hodes+ 1997] T. Hodes, R. Katz, E. Servan-Schreiber, and L. Rowe. Composable Ad Hoc Mobile Services for Universal Interaction. In *Proceedings of the Third ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM),* September 1997, 1–12.

[Kahn+ 1999] J.M. Kahn, R.H. Katz, and K.S.J. Pister. Mobile Networking for "Smart Dust." In *Proceedings of the Fifth ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM),* August 1999, 271–278.

[Leiner+ 1987] B.M. Leiner, D.L. Nielson, and F.A. Tobagi. Issues in Packet Radio Network Design. *Proceedings of the IEEE (Special Issue, Packet Radio Networks)* 75(1):6–20, January 1987.

[Malkin 1994] G. Malkin. RIP Version 2—Carrying Additional Information. RFC 1723 (draft standard). Internet Engineering Task Force, November 1994.

[Moy 1997] J. Moy. OSPF Version 2. RFC 2178 (draft standard). Internet Engineering Task Force, July 1997.

[Perkins 1996] C. Perkins. IP Mobility Support. RFC 2002 (proposed standard). Internet Engineering Task Force, October 1996.

[Postel 1981] J. Postel. Internet Protocol. RFC 791 (standard). Internet Engineering Task Force, September 1981.

[Prakash 1999] R. Prakash. Unidirectional Links Prove Costly in Wireless Ad Hoc Networks. In *Proceedings of the Third International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL M),* August 1999, 15–22.

[Rekhter+ 1993] Y. Rekhter and T. Li. An Architecture for IP Address Allocation with CIDR. RFC 1518 (proposed standard). Internet Engineering Task Force, September 1993.

[Royer+ 1999] E.M. Royer and C.-K. Toh. A Review of Current Routing Protocols for Ad-Hoc Mobile Networks. *IEEE Personal Communications* 6(2):46–55, April 1999.

[SDT 1995] A Survey of Defence Technology: The Software Revolution—"To Dissolve, to Disappear." *The Economist,* June 1995.

[Strater+ 1996] J. Strater and B. Wollman. *OSPF Modeling and Test Results and Recommendations.* Mitre technical report 96W0000017, Xerox Office Products Division, March 1996.

[Tsuchiya 1988] P.F. Tsuchiya. The Landmark Hierarchy: A New Hierarchy for Routing in Very Large Networks. In *Proceedings of ACM SIGCOMM '88,* August 1988, 35–42.

[Weiser 1993] M. Weiser. Some Computer Science Issues in Ubiquitous Computing. *Communications of the ACM* 36(7), July 1993.