

The Future of WLAN

MICHAEL W. RITTER,
MOBILITY NETWORK SYSTEMS

Since James Clerk Maxwell first mathematically described electromagnetic waves almost a century and a half ago, the world has seen steady progress toward using them in better and more varied ways. Voice has been the killer application for wireless for the past century. As performance in all areas of engineering has improved, wireless voice has migrated from a mass broadcast medium to a peer-to-peer medium. The ability to talk to anyone on the planet from anywhere on the planet has fundamentally altered the way society works and the speed with which it changes.

The changes triggered by wireless technology have only just begun. One of the most fundamental transformations in the last decade has been a result of the Internet, which provides us fundamentally (and, as yet, potentially) with instant access to all the information ever produced by the human race. It's easy to foresee the day when that ability resides in my pocket, accessible


anywhere, at any time. This capability will become a reality, however, only when wireless technology can provide ubiquitous high-speed data connections.

The most effective and elegant architecture for implementing this capability thus far is a wide-area mesh network that allows:

- Complete mobility of the end nodes
- Ubiquitous coverage without requiring expensive network management
- The ability to automatically replace the wired backbone with wireless links

WIRELESS IN THE LAST CENTURY

In the past quarter century, advances in wireless technology have unleashed a seemingly insatiable demand for mass-market peer-to-peer voice communication: Consider the citizen-band radio craze, the walkie-talkie market, the interconnection between radios and the public



Overcoming the
Top Ten Challenges in
wireless networking—
will it allow wide-area
mesh networks to
become ubiquitous?

The Future of WLAN



telephone system, and the incredible proliferation of cellular phones. This huge demand has precipitated a compelling impetus for improving wireless systems.

Wireless has also provided a host of new services just within the past two decades. First there was the transmission of exact time over radio waves, then the ability to determine where you were to within a few meters, anywhere on the globe, via the global positioning system (GPS). The list of wireless applications at my disposal each day is vast and varied. I control all my multimedia systems with wireless remote controls. I talk on my cordless phone, transmit music to my headphones wirelessly, and give presentations using a wireless microphone. I send music from my PC to my stereo system wirelessly. I connect to the Internet over a high-speed wireless connection in my house and at the local coffee shop. I get pictures from my front door to my television wirelessly. My front doorbell is wireless. I have a cellphone in my car, on my person, and for each of my family members. I receive my TV signals wirelessly from a satellite. I have radios in all my cars and in my house. My clock is synchronized to global time wirelessly. I have carried a two-way pager with me. My garage door and both of my cars unlock via a wireless "key." My computer mouse and keyboard connect to my PC wirelessly. We use wireless networks at work. I even eat most of my food after heating it up with microwaves. Wireless is everywhere, and getting more popular every day.

WIRELESS IN THE NEXT CENTURY

The biggest change in wireless over the past 10 years has been the availability of the unlicensed wireless spectrum. Before this, the evolution of wireless technologies was tied to a specific spectrum and specific protocols. Because of concerns about interference, the transmission of electromagnetic energy was highly regulated. This tended to channel technology into very conservative approaches.

During the past 20 years, however, other solutions have emerged. It became possible to connect a radio to a

computer that could control signals in new and complicated ways, at rates much greater than ever before. Combining this new application of computer technology with the unlicensed spectrum has dramatically accelerated the evolution of wireless technologies. In the past 10 years the data rate available over unlicensed wireless systems has changed by four orders of magnitude from 10 kilobits per second (kbps) to 100 megabits per second (Mbps). Nothing similar has ever occurred in the licensed spectrum in so short a time. What is different about the unlicensed spectrum?

UNLICENSED SPECTRUM CHANGES TECHNOLOGICAL EVOLUTION

Unlike the licensed spectrum, in which the FCC specifies the exact details of the radio's transmissions and inspects every deployment for compliance, the unlicensed spectrum allows you to build whatever you want as long as you follow the basic rules. The FCC has put little or no requirements on the services you can provide.

These factors make the unlicensed spectrum a perfect environment for producing a disruptive technology, which, in the case of wireless LAN (WLAN), it has. [See "The Innovator's Dilemma: When New Technology Causes Great Firms to Fail," by Clayton Christensen, *HarperBusiness*, May 2, 2000.] A disruptive technology is one that starts off as an inferior technology with less performance but at a cheaper price. The new technology improves faster than market needs evolve. Finally, the new, cheaper technology fulfills or surpasses all the market needs, and older, more expensive technology is displaced or disrupted. WLAN [refer to various IEEE 802.11 specifications at www.ieee.org] is such a disruptive technology. Today, it is replacing wired LANs and wired WANs; tomorrow it will be replacing other services, such as voice.

In many cases, providing a wireless link between buildings is cheaper than bringing fiber to the building. Digging a ditch for fiber typically costs \$100 to \$400 per foot.

A wireless link providing 100 Mbps over 300 feet can cost less than \$2,000. Longer ranges with higher data rates at minimal extra cost are available. In the coming decade, it will be cheaper to replace wires with wireless in every application. Table 1 compares data rates and costs.

An example of disruptive technology is taking place within the cellular telephone equipment market. Today, cellular systems have complicated definitions that create numerous dependencies up and down the protocol stack. Thus, they are tailored for a vertically integrated (i.e., proprietary) solution. Vertical markets tend to emerge early in the development of new technologies. For example, the computer industry was once vertically integrated. You would buy a mainframe with built-in disks, software, programming languages, input terminals, special memory, proprietary networks, etc.—all from the same vendor. Today, the vast majority of processing power is in personal computers that are made by hundreds of manufacturers, with components that are essentially interchangeable.

The computer industry has moved from a vertical market, where one company builds the entire system, to a horizontal market in which several suppliers of subsystems compete directly with one another while systems integrators assemble the components. As competitors enter the market, standards begin to emerge that allow groups of smaller companies to build solution components that interoperate with one another. This leads to the erosion of the vertical model into a horizontal one, where serious competition exists at every level of the system.

WLAN is poised to spearhead the transition of the entire cellular industry into a horizontal market. WLAN is a completely horizontal market, with standards in place. The price of WLAN technology has fallen at almost 40 percent per year for the past 10 years, and there is no reason to believe this will slow down.

As computers, radios, and networks become smarter:

- Handoffs will become faster.
- Quality of Service (QoS) problems will be solved.
- Backhaul will become cheaper (as it becomes entirely wireless).

TABLE 1

Technology	Data Rate (kbps)	\$\$/kbps
GPRS	53	\$118.91
1xRTT	153	\$25.42
Ricochet	248	\$2.43
WLAN	5,500	\$0.27

WLAN technology is a disruptive data technology in both performance and cost. The first column lists the maximum user data rates available, the second lists cost, assuming blanket coverage of a large area. [Capacity and cost comparison of Ricochet with General Packet Radio Service (GPRS) and 1xRTT by LCC International Inc., August 2, 2001; WLAN by personal estimates.]

- The small cell sizes of WLAN will become less and less of a barrier.

As this transition begins to accelerate, the current cellular infrastructure will be replaced with faster, cheaper, standard WLAN equipment.

Unlicensed radio technologies, such as WLAN, are changing many industries. Soon, every portable device will be wireless, as the cost of adding wireless connectivity will be nil. Eventually, wireless will be available everywhere. The devices will form a ubiquitous mesh network, and broadband wireless data access will be the norm.

To simplify wireless deployment, a mesh architecture is the most likely approach to succeed. It allows wireless nodes to connect to each other automatically using the most efficient path to move data between the nodes. Although this architecture is the most flexible, it is the most difficult to design and deploy. The ubiquitous wireless world I describe could become a reality over the next 20 years, but only if several challenges are overcome.

TOP TEN CHALLENGES IN WIRELESS DATA

Let's assume that we want to build a very high-speed, ubiquitous wireless data network that can be sold for a fixed monthly fee costing about the same as any other broadband method. First, we must meet what I call, in David Letterman fashion, the Top Ten Challenges in Wireless Data:

10. Eliminate outdated systems that tie up spectrum (broadcast television and radio, any analog system, any system that is not spread spectrum).

9. Provide the fast handoffs that will be required for continuous mobile connectivity, as cell sizes will continue to decrease.

The Future of WLAN



8. Solve the hidden-terminal problem, which is really a question of coordinating a large number of radios to reduce interference.
7. Coordinate individual radios so that Quality of Service can be guaranteed in a mesh network.
6. Reduce power consumption of the entire system, especially user devices.
5. Create standard ad hoc routing and MAC layers that work for large meshed networks of mobile nodes with high throughput and low delay over many hops.
4. Provide cheap, smart antennae and the protocols that go with them. Without directional antennae, interference problems become exponentially worse.
3. Provide a cheap, wired backbone to enable inexpensive connectivity to the wireless mesh.
2. Develop algorithms for maximizing system throughput and capacity in large meshed networks.
1. Integrate unlicensed wireless securely and transparently into existing networking systems, such as wired enterprise Ethernets, the cellular system, and the public switched telephone network.

I will discuss each of these challenges in turn, direct readers to the current research on each topic, point out likely approaches to solve these problems, and discuss the advances required to make these approaches viable.

10. Eliminating inefficient usage of spectrum. This is one of the most difficult problems to solve, because constituencies that have monopolies on using the spectrum benefit greatly from them and will not easily give them up. This problem has only two solutions, and both must be applied with care. The first is slowly and steadily to force technologies to change by providing a combination of the right incentives and the right threats. The FCC has taken this approach, attempting to redefine the type of usage allocated to many areas of spectrum.

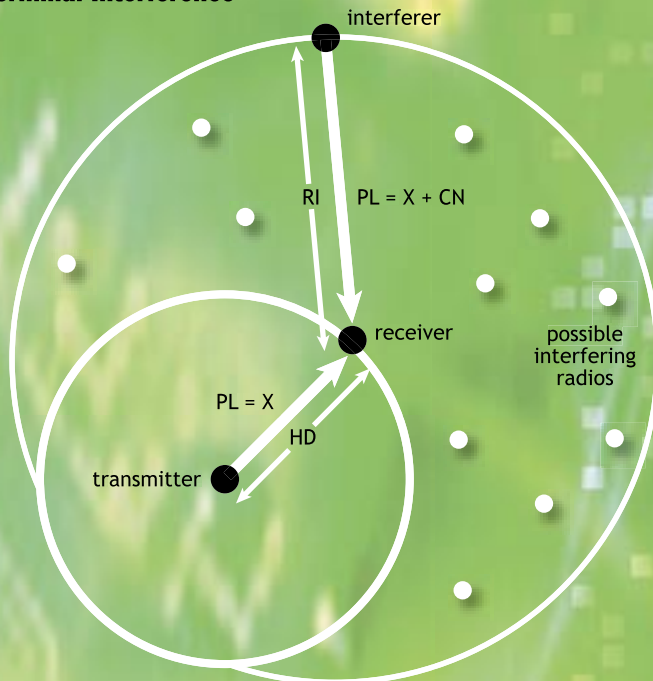
This reallocation of spectrum is slow, with the “chipping away” of the inefficient usage of spectrum proceeding at a rate I estimate to be less than one-fifth

of 1 percent per year. At this rate, it will take more than three centuries to free up half of the spectrum suitable for ubiquitous-coverage mesh networks. Fortunately, this process is beginning to accelerate as a result of the hard work of many dedicated people. (For example, on Capitol Hill, Rep. Fred Upton (R-Mich.) introduced H.R. 1320, the Commercial Enhancement Act, which would create a fund allowing government agencies located on commercially attractive spectrum to relocate to other parts of the spectrum. Auction proceeds from the sale of vacated spectrum would be used to cover relocation costs.) In the long run, these efforts will have a huge effect on communication capacity for individuals.

The second solution is a technical one, exemplified today by new standards that the FCC is defining for new types of spread spectrum radios, particularly the new rules covering ultra-wideband (UWB) radios. [For more information, refer to the FCC’s Ultra-wideband Radio Specifications, ET Docket 98-153, First Report and Order, adopted Feb. 14, 2002, released April 22, 2002.] These radios spread their signals over many gigahertz of spectrum, at such a low power level and duty rate that other users of this spectrum are unaware they are transmitting. Other modulation techniques allowing multiple uses of the same spectrum (with manageable interference) are in various stages of development. Producing these types of radios, and convincing the FCC to let them share spectrum, is a crucial task for those who wish to provide ubiquitous broadband wireless coverage.

9. Fast handoffs. One of the greatest system design concepts in the evolution of wireless technology is the notion of cellular deployment and the huge reuse of spectrum it provides. Reusing spectrum every few miles or every few hundred feet provides an essentially unlimited capacity for wireless communications. The smaller the cell size, the larger the communication capacity per unit area [see “The Capacity of Wireless Networks,” by P. Gupta and P.R. Kumar, *IEEE Transactions on Information Theory*, pp. 388-404, Vol. IT-46, No. 2, March 2000]. This

Hidden-Terminal Interference



PL = path loss in dB of the radio signal, CN = the ratio of the carrier power in the radio signal to the noise in dB required for successful reception of a signal, HD = hop distance from the transmitting radio to the receiving radio, RI = radius of interference, or the distance within which transmitting radios can interfere with the transmission between the two illustrated radios. Shown and typical, CN is > 0 ; thus, $RI > HD$. This implies that radios that can't directly talk to one another can interfere with each other.

feasible. Standards bodies today are defining the interfaces required to make this possible [see 802.11 IEEE Task Group k and f, protocol handoff definitions and resource definitions, at www.ieee.org].

8. Hidden-terminal problem. This problem is particularly difficult to overcome for a large group of meshed radios. It occurs because a radio needs to receive a signal that is above the noise level by a certain amount before it can decode it correctly. To maximize the capacity of a mesh network, you never want a packet to fail. When a “hidden” terminal transmits a packet without coordinating with other nearby radios, it can easily interfere with another conversation going on nearby. (See Figure 1.)

As the load on the network

imposes huge pressure for very high-speed, short-distance radios. Thus, as cell size naturally shrinks, the available data rate increases.

This is exemplified by the large deployment of WLAN in homes and enterprises. This large market will drive down the cost of deploying cellular-type systems. Today the market is rapidly approaching the point where it is less expensive to put a WLAN node on every streetlight than to deploy a conventional cellular system to cover the same area (especially when taking into account wireless backhaul).


So why isn't this done today? The main limiting factor is the ability to provide fast handoffs over cells sized a few hundred feet in diameter. Typical vehicle speeds of a hundred feet per second mean switching cells every few seconds. This seems impossible to contemplate today, but think about the hundreds of millions of decisions a 2-GHz processor can make in a fraction of a second. The ability to find the radio in the next cell and to hand off to it in this same short period of time seems eminently

increases, interference increases, and the corresponding efficiency of the network decreases, eventually driving the throughput of the network to zero as everyone tries to transmit all the time.

There are three potential ways to fix this problem, and all of them are difficult. Advances in processing power will enable implementation in the near future, however.

- The first solution is to coordinate all the radios' transmissions. The problem with this is that radios can interfere with each other but may be incapable of talking to each other; therefore, it may be impossible for them to directly coordinate with each other. The solution is to coordinate with their distant neighbors over the network. In cellular networks this is done by programming in different frequencies to be used by the radios in different positions. New algorithms need to be invented to support mobile mesh networks. This increases the level of complexity, because the radios must know where they physically reside before they can perform this coord-

The Future of WLAN



dination. (Of course, this can be determined wirelessly using a system such as GPS—one wireless technology helping another.)

- Another way to address the hidden-terminal problem is to change the carrier-to-noise ratio (CN) to a negative value. This simplifies one dimension of the problem because the hop distance (HD) can now be larger than $2 \cdot RI$, and radios can successfully talk to each other from farther away than they can interfere with one another. This simplifies coordination, but does not entirely solve the problem. Although difficult to implement with spread spectrum radios and advanced signal processing, this reduction in CN is possible. Nothing is free, however. When CN is negative, the next layer of the onion is the so-called “near-far” problem. To decode a signal with a negative CN, all signals must arrive at the receiver with approximately the same power level so that N, the noise generated by the other radio signals, does not become too large. One large nearby signal can drown out all the small signals from distant radios. Some form of coordination still must be implemented. The best way to perform this coordination has yet to be determined.
- The final proposed solution takes into account the fact that N is not really noise; it is the sum of the signals of all other radios. Advanced signal-processing algorithms can take advantage of this and effectively mask the competing signals, thereby greatly reducing N and effectively increasing CN to reduce the chance of interference. This is not as easy as it sounds (not that it sounds that easy). The received signals must be processed once for each interfering signal, as one cannot tell *a priori* which signal is destined for the receiver. With exponential improvements in signal-processing power, this problem can be solved in the coming years. It can be further simplified by supporting smart antennae at each radio. That, however, is a separate research topic of its own, to be discussed later in this article.

7. Coordinating QoS. I used to promote the idea of providing Quality of Service by increasing the available bandwidth. That solution works well for wired connections, where you can easily increase the additional bandwidth or add a link that won't interfere with the original one. For wireless, however, you want to make the most efficient possible use of the airwaves because you can't just magically add more links between two radios to increase bandwidth. The research that has already been done using asynchronous transfer mode (ATM) to determine how to provide QoS can now be applied to the real problem of making wireless links more efficient. The truly difficult problem, which has only begun to be addressed, is how to provide QoS in a meshed network over multiple links between multiple radios. Although some preliminary work has been done in this area [refer to “A Routing Architecture for Mobile Integrated Services Networks,” by S. Murthy and J.J. Garcia-Luna-Aceves, *ACM Mobile Networks and Applications Journal*, 1998], practical solutions have yet to be produced. The additional level of complexity makes this an interesting problem to solve.

6. Power consumption. The need to reduce power consumption is one of the most challenging and interesting topics in wireless engineering, which not only tackles theoretical problems, but also requires complicated physical solutions before any real system can be implemented. The designer is always trying to push the entire system to its theoretical limits. Today's cellphone solutions have made huge strides in conserving battery power. This has been achieved with a relatively fixed throughput. The next challenge is to keep this trend intact while increasing the throughput by several orders of magnitude.

One method of decreasing power consumption is to bring the radios closer together or, in other words, to create a ubiquitous mesh network. To bring radios closer together, they must become exponentially cheaper, as the number required increases (at best) by the square of the ratio of decrease in their average separation. Thus, using

a mesh architecture, you can decrease the power consumption of a radio system simply by making the radios cheaper.

There is also a great push to increase the power density of batteries or other power-storage devices for reasons of portability. Wireless protocols and system design should always strive to minimize power consumption. It is one of the key criteria when building a system. In a wired system, this requirement is usually ignored because it has little effect on the performance of the system. For a wireless mobile system, carefully managing this specification is crucial to making a usable system.

5. Designing mesh network protocols. Designing efficient mesh networking protocols is critical. The industry still has a long way to go to standardize on anything approaching an efficient routing protocol for mesh networks. The simplest way to approach mesh networks today would be to implement a meshed routing protocol on top of existing WLAN protocols. Progress has been made in this direction [see www.ietf.org/html.charters/manet-charter.html]. Without some modifications at the MAC or physical layer, however, this approach is doomed to system inefficiencies and possible total collapse in large networks, because of MAC layer problems [refer to “Does the IEEE 802.11 MAC Protocol Work Well in Multihop Wireless Ad Hoc Networks?” by Shugong Xu and Tarek Saadawi, *IEEE Communications Magazine*, 2001].

The MAC and routing layers must be designed in conjunction with each other to wring out the last bits of efficiency. If it is not done carefully, a design that works eminently well in a WLAN will collapse in a large meshed network. To produce the ideal solution, the physical layer should also be designed from scratch. Individually maximizing the design of each layer is not sufficient. An efficient design must include maximized system throughput. These challenges make the design work a difficult, albeit interesting, problem that has yet to be solved. Pioneering work has been done [refer to “Improving TCP Performance over Wireless Networks at the Link Layer,” by C. Parsa and J.J. Garcia-Luna-Aceves, *ACM Mobile Networks and Applications Journal*, Vol. 5, No. 1, 2000, pp. 57-71], but many fundamental problems in this area are still not understood.

4. Smart antennae. This is one of the key areas that wireless engineers have yet to put into common use, especially in the area of mesh network design, where it would be extremely valuable [see “MicroCellular Data Network (MCDN): Performance and Capacity of a Broad-

band Mobile Wireless Technology,” by Robert Friday, Michael Ritter, and Arty Srivastava, *Network-Interop 2000*, Las Vegas]. This problem becomes extremely challenging when the radios in the mesh network are moving. The research in this area has been excellent, and the theoretical limits of the technology in stationary situations are well understood [refer to “On Limits of Wireless Communications in a Fading Environment when Using Multiple Antennas,” by G.J. Foschini and M.J. Gans, *Wireless Personal Communications 6*: pp. 311–335, 1998, Kluwer Academic Publishers]. Achieving a practical implementation that comes anywhere near the theoretical limit in a mobile ad hoc environment, however, is largely uncharted territory.


With increases in signal-processing power, approaching the theoretical limit at a *reasonable* cost should be possible in the near future. Knowing the theoretical limits may be intellectually satisfying but is insufficient in a commercial environment unless you can implement a design that produces these results in a commercially viable configuration. This task is exponentially more difficult and, of course, more satisfying and rewarding.

3. Cheap wired backbone. You would think that with the surplus of fiber supposedly installed in the world, connecting a wireless meshed network of radios to the wired backbone would be no problem. Wrong. The apparent glut of fiber is really a glut of bandwidth between specific points and an absolute dearth of high-speed connectivity everywhere else. Meshed wireless networks can help solve this problem because they can convey the signals to where the fiber exists. Because the fiber is hardly anywhere (less than 1 percent of commercial buildings have fiber on premises today and less than 5 percent of buildings have fiber that actually passes by their front doors), a cheap, wired backbone that is also relatively ubiquitous would create a significant savings.

A huge unmet need exists for high-speed wired connections, and wireless networks will suffer because of it. One of this nation’s highest priorities should be to resolve this problem quickly. In return for the monopolies the telephone companies have, it seems a small price to pay to require them to provide this service. Today, they are required to provide ISDN connectivity anywhere in their covered areas.

The phone companies cannot be entirely blamed, however. Competition has not brought about the desired results. Cable companies have huge amounts of bandwidth available over their infrastructure but have been able to provide only small data pipes for their customers. Something is wrong with the incentives in the system.

The Future of WLAN



This is probably the most difficult problem on the entire list, because it requires changes in both the government and the free market. The FCC and Congress should make this one of their highest priorities.

2. Maximizing system throughput in a meshed wireless network. This is the key problem that must be solved to provide high-speed ubiquitous wireless coverage in an efficient manner. It essentially takes all of the preceding problems and puts them together to produce the best and most efficient system. This is an interesting problem because the theoretical upper bounds have yet to be convincingly determined [see “Position Based CDMA with Multiuser Detection (P-CDMA/MUD) for Wireless Ad Hoc Networks,” by Teresa H. Meng and Vokan Rodoplu, IEEE 5th International Symposium on Spread-Spectrum Technology and Applications, NJIT, New Jersey, Sept. 6-8, 2000; and Decentralized Channel Management in Scalable Multihop Spread-Spectrum Packet Radio Networks, by Timothy Jason Shepard, Ph.D. thesis, MIT, 1995, theses.mit.edu:80/Dienst/UI/2.0/Describe/0018.mit.theses%2f1990-125?abstract=]. Even to approach a reasonable estimate, we must make many simplifying assumptions. My guess is that the theoretical limits defining this problem will be determined within the next few years, but that practical solutions approaching these limits are still a decade away.

1. Integration of wireless data into existing systems. If you solve the first nine problems, you are left with this one: The wireless meshed network must intercommunicate with all the existing legacy systems out there. The interconnections need to be efficient and economically practical. The meshed network can't require the legacy systems to be extensively modified. Several unexpected problems have come up. For example, TCP/IP was originally designed to run over lossy, highly variable delay networks and still deliver reasonable performance. The biggest problem for most users of TCP/IP, however, turned out to be that it was not efficient under

the load of multiple users. The design was changed to make it efficient under load [see “Congestion Avoidance and Control,” by V. Jacobson, *Procedures, SIGCOMM*, Stanford, CA, Vol. 18, No. 4, August 1988]. Now it doesn't work well over lossy networks with highly variable delays. There are two avenues of attack for these types of problems, and both should be followed through in parallel:

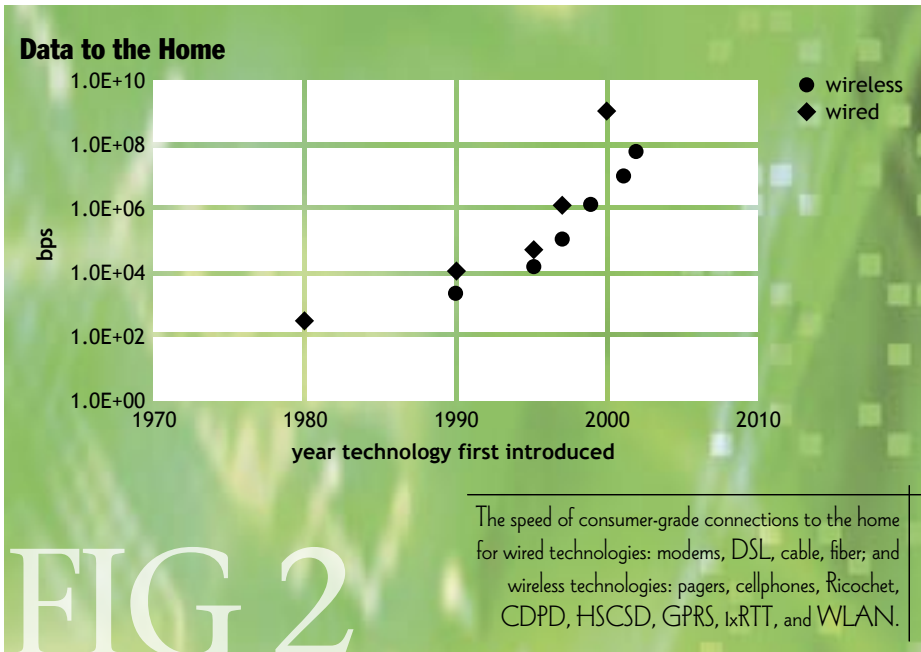
- Modify the existing protocols to be more flexible.
- Make the wireless mesh network look more like a wired network with predictable delays and minimal losses.

Several other systems integration opportunities, ripe for solutions, present themselves. For example, voice over IP (VoIP) over WLAN and integration of this feature transparently into the existing cellular systems is a problem within reach of a solution. Integration of WLAN systems into existing operators' service portfolios is also a difficult problem that is spawning its own industry [for example, Mobility Network Systems' current product offerings: www.mobilitynetworks.com; and competitors' offerings: www.nokia.com, www.cisco.com, www.transat-tech.com, www.adjungonet.com, www.tatara.com, etc.], whose designs are rapidly becoming standardized [for example, Third Generation Partnership Project (3GPP) efforts in such groups as SA1 and SA2: www.3gpp.org; and work in the Wireless LAN Taskforce of the General System Mobile Association (GSMA) www.gsma.org].

The wireless data industry has many challenges ahead. Advances in other fields are beginning to provide the necessary foundations for overcoming them in the next few years. The problems presented by unlicensed wireless data are some of the most challenging in the field.

THE FUTURE OF WIRELESS

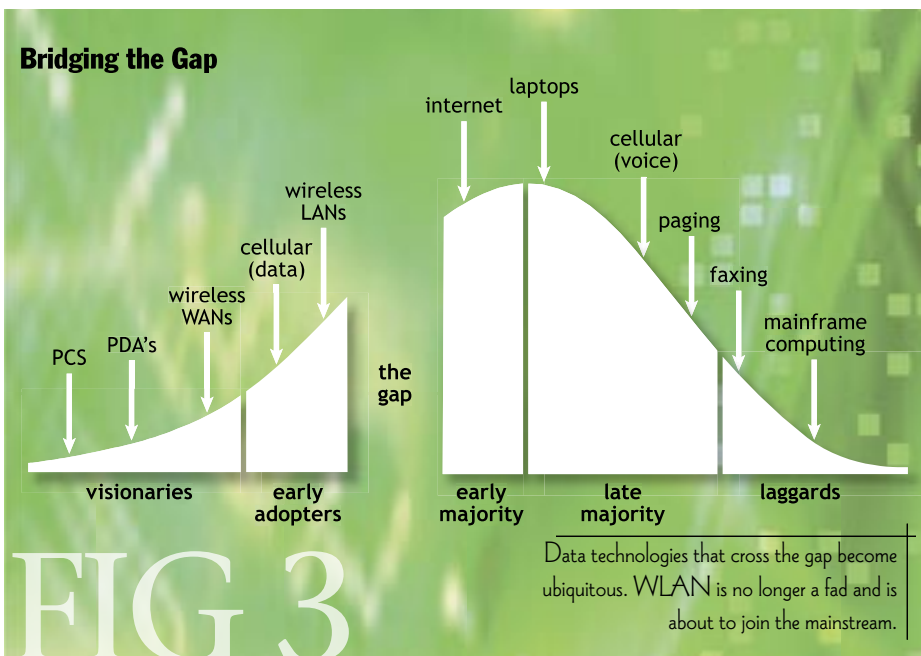
Networking speeds available in the consumer space have continued to increase over the years, as shown in Table 1, which extrapolates this trend into the future at the same rate and shows a doubling period of just over 18 months. There is no reason to expect this to change in the near future. Development of standards to support the data rates in Figure 2 is under way today—i.e., the 802.11a, b,



If the market penetration crosses the gap, it becomes a ubiquitous technology; if it doesn't cross, it becomes just a fad and slowly fades away. WLAN is poised to leap across that gap.

During the next two decades, wide-area mesh networks will face this same obstacle. If the problems described here are solved, this technology will also be capable of crossing the gap and becoming a ubiquitous technology. Q

MICHAEL RITTER is chief technology officer at Mobility Networks. Prior to that, he was CTO at Metricom, where he was responsible for architecting and developing a wireless network and infrastructure for the nationwide Ricochet Wireless Internet Service. Ritter also was manager of open systems development at Apple Computer, where he was responsible for MacTCP. Earlier, he was a manager at Sun Microsystems. He has held consulting and research positions at Lockheed Space and Missiles Company, Stanford University, Lawrence Livermore Laboratories, and was a visiting staff member at Los Alamos National Laboratories. Ritter holds a bachelor's degree in physics from Stanford and an MS, MA, and Ph.D. in physics from Yale. He is a member of IEEE and ACM. E-mail: mike.ritter@mobilitynetworks.com.



g, and 60-GHz proposals [refer to IEEE task groups and proposals, www.ieee.org].

So where are we today? Figure 3 is a chart that plots various technologies against their distribution in the general population [see "Crossing the Chasm," by Geoffrey Moore, *HarperBusiness*, revised edition, August 20, 2002]. The gap between early adopters and early majority is the defining moment for a particular technology.

ratories. Ritter holds a bachelor's degree in physics from Stanford and an MS, MA, and Ph.D. in physics from Yale. He is a member of IEEE and ACM. E-mail: mike.ritter@mobilitynetworks.com.

The author would like to thank Mike Pettus, Francie Miller, Stan Potts, Judy Lindo, and Robert Friday for contributions to this article.

© 2003 ACM 1542-7730/03/0500 \$5.00