Advanced Wireless Technologies:
MIMO Comes of Age

A Farpoint Group White Paper

Document FPG 2003-242.2
December 2003

Farpoint Group

If there is one great truth in data networking of any form, it's that we never seem to have enough performance. Effective throughput requirements only increase over time. Anyone who remembers the days of 9.6 Kbps modems and 10 Mbps Ethernet likely appreciates the difference between those venerable technologies and the networking capabilities in place today. Indeed, the probability that anyone would voluntarily give up a cable modem or xDSL connection to the Internet is so small as to be statistically insignificant.

While not yet obvious to everyone, so it also is with wireless LANs. Anyone who remembers the days of one and two Mbps WLANs has zero desire to return to them. And even with today's 11 and 54 Mbps technologies, the need for greater throughput is ever present. If one can synchronize one's notebook in ten minutes, wouldn't five minutes be better – and one minute even more desirable, especially when one needs to get to the airport right away? Add to this the requirement for time-bounded communications (for voice and multimedia), and the need for ever-greater network capacity, and performance remains a key challenge for WLANs.

While, in truth, we will likely never have enough performance in any network, again given growing application requirements and user expectations, the provision of higher throughput is especially challenging in wireless LANs. The reason for this is simple – the radio channel, that property of the universe utilized to move information over the air, has a *highly-variable* nature. Unlike the relatively stable environment that exists on wire, cable, or fiber, the ability of the air to carry information can and does change over time, and often from moment to moment. Given this fundamental variability and the overhead inherent in any networking protocol, the actual throughput available from a 54 Mbps connection is often less – *and usually much less* – than this peak number. As a consequence, we need to improve the performance of wireless LANs at the *physical layer* if we are to achieve higher throughput in the future. One popular (if non-standard) approach to date has been to gang together multiple radio channels (sometimes called "channel bonding"), for example yielding 108 Mbps via the simultaneous utilization of two 54-Mbps channels. We can also use compression and related techniques to gain some additional advantage in many cases. But the ideal solution would be to come up with a technology that simply packs more information per unit of bandwidth and time (see the sidebar *Understanding Spectral Efficiency*).

The classic technique applied in this situation in to improve what is known as *modulation efficiency* – the number of bits per unit of bandwidth and time we can reliably fit onto the air in any given situation. This approach turns out to have some fundamental physical limitations, again due to the variable nature of the radio channel. Radio signals are subject to serious degradation as they move through space, primarily due to the distance between transmitter and receiver, interaction with objects in the environment, and interference from other radio signals and reflections of the signal in question itself (known as *multipath*). All of these artifacts result in a number of forms of *fading*, the loss in power of a radio signal as it moves from transmitter to receiver. Eventually, signals become so weak that they can't be reliably received – there simply isn't enough signal to demodulate (and sometimes even detect) the transmission and extract the communicated information from the carrier. Thus, more aggressive modulation makes sense at shorter distances, and modern radio protocols (including most 802.11 PHYs) allow for the modulation to become more or less aggressive as distance and prevailing radio conditions so dictate - which is why 802.11a and .11 g vary in speed from six to 54 Mbps, and .11b from one to 11 Mbps. Regardless, there are fundamental limits to reliability as modulation efficiency improves. While there is some possibility that advances in conventional signal processing technology will at least partially overcome this problem at some point in the future, a far more reliable and productive technique is available today, and has recently been put into practice in a production wireless LAN. It's called *MIMO – multiple input, multiple output*. While we'll explore MIMO in some detail in this document, for the moment the best way to explain MIMO is as the addition of another dimension in the radio channel – a *spatial* dimension – allowing a more complex but inherently more reliable radio signal to be communicated. Greater reliability means greater throughput. By analogy, MIMO is to radio as 3D computer graphics is to 2D. The richness of today's video games, for example, is due to their inherently 3D nature. Anyone who's experienced the graphics of a modern game console wouldn't go back to the 2D world of *Pong* and *Space Invaders* except perhaps for an occasional chuckle.

MIMO – Definition and Properties

First, by way of explanation, the "input" and "output" in MIMO refer to how we access the radio channel itself. (see Figure 1). Whereas a conventional radio (which is known as SISO, for single-input, single output) will have one input (a transmit antenna) and one output (a receive antenna), a true MIMO system uses at least two transmit antennas, working simultaneously in the service of a single logical channel, and at least two receive antennas at the other end of the intended connection. In reality, the number of receive antennas in a MIMO system will usually be greater than the number of transmit antennas, and performance generally improves with the addition of more receive antennas. Just going from a single antenna to two MIMO antennas can result in a 10 dB (10X) improvement in signal-to-noise ratio (SNR), a key measurement of reliability. Adding a third antenna adds an additional five dB, an almost 4X additional improvement.

Understanding Spectral Efficiency

The key metric for measuring how efficiently the airwaves are being used is *bits per second per Hertz*, or *bps/Hz*. One Hertz is one cycle per second, so the question is how many user data bits can be packed into a single Hertz every second. Note that the more complex the modulation scheme, the more efficient the bandwidth utilization. But at some point the modulation gets so complex that it's either too expensive to build (and power) the circuitry required to produce a more desirable result, or the resulting waveform is simply too difficult to demodulate after it has been through the wringer of the radio channel itself.

Consider, as an example, the OFDM scheme used in 802.11a and 802.11g, which yields (at its maximum rate) 54 Mbps in a 20 MHz. channel. This gives us 2.7 bps/Hz. Note this is the best possible speed and doesn't take into account the variable performance of the standard (which actually ranges from six to 54 Mbps, and varies with range, channel conditions, and other factors), nor other overhead related to the 802.11 MAC and network-layer (and above) protocols such as TCP/IP. MIMO techniques currently allow raw throughput to rise to 108 Mbps, which doubles spectral efficiency to 5.4 bps/Hz. Coupled with the improved range that MIMO can provide, and the net benefit of going with a MIMO solution can be more than double that of the traditional approach.

As we noted above, a radio channel is defined by a certain amount of *bandwidth* available between two given *frequencies*, usually called a *band*. The band in use is most often determined by *regulation*; that is, the set of rules relating to spectrum allocation in any given political jurisdiction, typically a country. Further, the nature of a specific band is additionally determined by the concept of a radio *channel*, which is usually much smaller than the band.

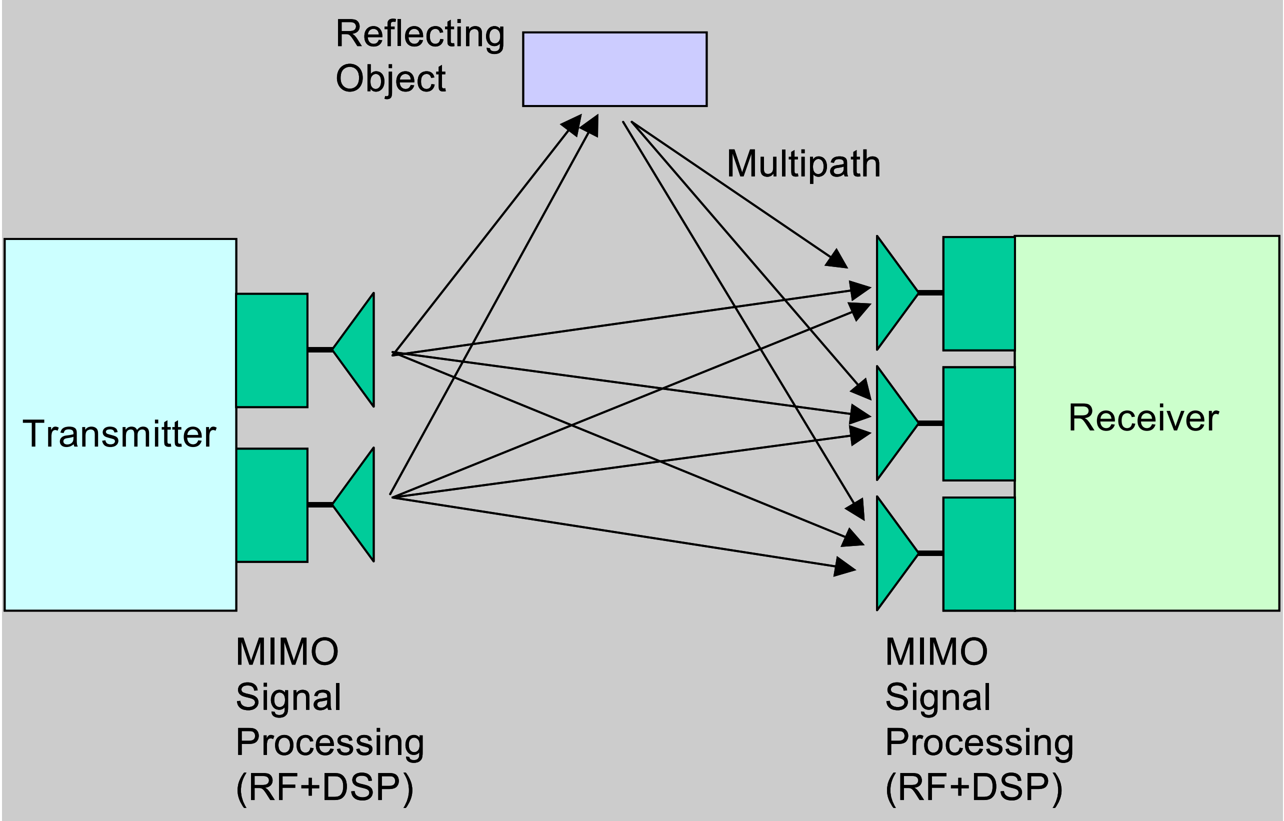

Figure 1: MIMO uses multiple transmit and receive antennas, and depends upon interactions with the environment in the form of multipath for its benefits—a counterintuitive element in the technology. *Source:* Farpoint Group

For example, in the 2.4 GHz. unlicensed band, the IEEE 802.11 standard defines 11 channels in the United States of 20 MHz. each, three of which (channels 1, 6 and 11) are non-overlapping. A final complication is that radio propagation is by its very nature *non-deterministic*; that is, it is impossible to tell exactly how (or even if) a given radio signal will travel from the transmitter to the receiver and arrive in a form that allows detection and demodulation to take place. In addition to these environmental concerns, the propagation of radio waves through space also varies with *frequency* – each band has its own set of propagation characteristics. For example, some frequencies just naturally travel through walls and other solid objects more reliably than others.

All of this makes the role of the antennas used at both ends of the connection critical. Antennas are often marginalized as to their contribution in successful radio connections, and many antennas, especially those in wireless LANs, are often no more than simple squares of copper. But antennas have perhaps the most critical function of all in determining the success of a given radio link – their role is not unlike that of tires on a car. Tires, after all, are the only part of the car that actually touches the road (in normal operation, anyway!). Proper tire pressure and tire type and configuration for a particular car can make a dramatic difference in handling, fuel economy, and response to emergency and other unusual or marginal – but nonetheless critical – conditions.

It stands to reason, then, that if we can improve the performance of the antennas used in a WLAN, we can significantly affect the overall performance of the radio and the radio link itself - and that is exactly what MIMO is designed to do. MIMO introduces a third *spatial* dimension beyond the frequency and time domains which otherwise define the radio channel. A radio signal from a point-source (single) antenna will typically bounce around quite a bit during transmission, particularly indoors, as it interacts with objects in the environment. A key result of these interactions is *multipath fading*, as the signal interferes, often destruc-

An Overview of Multiple Antenna Technologies

While it's easy to argue that the antenna is the most important part of the radio, it's a lot more complex to argue the features and benefits of differing approaches to antenna implementations, if for no other reason than the mathematics becomes intractable for all but the most proficient engineers. In principle, a receiving antenna behaves a bit like a magnet – it's designed ("tuned") to a particular signal at a given frequency. Antennas can actually be designed to handle a broad range of frequencies, power levels, and can also be designed to at least partially reject interference and signals coming from a direction other than that intended (or, correspondingly, to focus energy in a particular direction when transmitting). Let's examine some of the common approaches used in antenna design.

Almost all wireless LAN antennas incorporate *switched antenna diversity* in the receiver. This relatively simple technique uses two receive antennas; if one antenna is in a deep fade with respect to the transmitter, there is some statistical likelihood that the other will not be. Thus, a diversity receiver system allows the electronic part of the receiver to be switched between the two antennas based on signal strength – whichever antenna seeing the strongest signal will be used at any given moment in time. Note, however, that there's no additional signal processing in this case, just switching, so while reliability is enhanced there are no fundamental improvements in the signal itself involved here.

As we noted above, directional antennas can be used to focus energy in a given direction, or reject signals coming from other than the desired direction. These can be quite useful when at least one end of a connection is fixed, and the radio knows the location of the other end of the connection so that the signal can be aimed in the proper direction. In a wireless LAN system, APs are almost always fixed. However, the mobile stations do not know where the APs are, and the clients typically move over time, so directional antennas cannot be effectively used on the client end of the connection and only in certain cases at the AP. While directional and even sectorized coverage is possible at the AP, the effect of multipath and other environmental interactions mitigate its usefulness to a great degree.

One approach to overcoming this problem is the use of "intelligent" or "smart" electronic, multi-element antennas usually based on phased-array technology. A phased array is usually implemented as a (rather large) flat panel with a relatively large number of active antenna elements. The elements can be electronically steered so as to focus energy from the antenna in a given direction, and the steering can even be dynamic so as to allow for some degree of tracking of a mobile client. Unfortunately, phased arrays are certainly not mobile, at least somewhat obtrusive, and expensive. And while they can be valuable in wide-area applications, their use indoors is questionable.

A number of what we might call "minimalist" multiple-antenna approaches, known as single-input, multiple-out (SIMO) or, conversely, multiple-input, single-output (MISO) systems can also provide improvements in performance, but a full MIMO implementation is clearly the optimal approach. MIMO techniques provide very high levels of performance at costs more akin to simple diversity techniques. Most importantly, MIMO provides very high performance without exceeding regulatory limitations on power or bandwidth, and without the need for aggressive modulation techniques that to date have proven unworkable especially in low-cost, battery-powered, highly-mobile implementations.

tively, with itself. MIMO actually *takes advantage* of multipath, using signal processing implemented on digital signal processor chips (DSPs) with clever algorithms at both ends of a MIMO channel. Somewhat counterintuitively, MIMO actually *depends* upon multipath to function correctly, making it even better suited to in-building applications like WLANs.

A core challenge in MIMO, beyond understanding the theoretical and practical complexity of building such a system in the first place, is how to squeeze a MIMO implementation into a form factor that is suitable for a wireless LAN. This means not only physical size, but equally importantly managing power consumption in a process that is inherently computationally-intensive, has multiple transmitters and receivers, and must operate in battery-powered devices. Interestingly, though, MIMO-based WLANs are now available, and we expect their role to be vital to the evolution of future WLAN systems – perhaps the single most important technical contribution to the future of the entire WLAN industry.

MIMO in Wireless LANs

To date, the evolution of wireless LANs has been driven by a three key factors, as follows:

- *The 802.11 Standard* – The release of the original 802.11 standard in 1997 legitimized WLAN technology in the minds of enterprise and residential users alike. While other standards and proprietary approaches alike attracted attention for a while,

A Brief History of MIMO

The first applications of MIMO were in point-to-point, line-of-sight microwave links, which remain popular today but were one of *the* major applications of radio in the late ‘70s. Given the limited spectrum available at the time, the need for more-efficient bandwidth utilization was a critical requirement in building higher-throughput radios. Among the vendors experimenting with MIMO and producing products were Bosch, Harris, and Siemens. Performance as high as 10 Mbps in a 2.5 MHz. channel was reported, a remarkable achievement for the time. Early MIMO systems used two receive antennas and both vertical and horizontal signal polarization – a technique that aligns radio waves accordingly. Each antenna would be able to receive both the vertical and horizontal signals, and a “space-time equalizer” was then used to cancel otherwise interfering signals and thus improve reliability. These systems were extremely large and expensive and were not suited for consumer or indoor applications.

In 1996, work done at Stanford University by Greg Raleigh and VK Jones (both now at Airgo Networks) mathematically proved multipath to be *essential* to realizing the full benefits of MIMO, a counterintuitive, as noted elsewhere in this document, but critical conclusion. Raleigh and Jones were the first to prove that with multipath the capacity and spectral efficiency of a MIMO system can be increased indefinitely.

The next highly-visible leap in MIMO came in the 1998, when Bell Labs announced their BLAST (Bell Labs Layered Space-Time) technology. The BLAST prototype used eight transmit and 12 receive antennas, and resulted in reported performance of 20 – 40 bits/second/Hz. in the lab. While Lucent has yet to commercialize the technology, BLAST raised the level of awareness of MIMO dramatically. The most exciting development in MIMO to date, however, is its application to the rapidly-growing field of wireless LANs.

Airgo Networks’ MIMO PC Card. Note multiple antennas on right. *Source:* Airgo Networks

In 2003, Airgo Networks announced the first commercial 802.11-compliant MIMO implementation in their AGN100-series products (photo at right). While remaining fully backwards-compatible with existing 802.11 products, the AGN100 is the first WLAN product to yield 108 Mbps in a single radio channel, and is indicative, we believe, of the direction that will be taken by the IEEE 802.11n Task Group, now developing the next-generation physical-layer (PHY) standard for wireless LANs.

802.11 is today the only wireless LAN standard that matters. Note that the 802.15.3a wireless personal-area network (WPAN) standard now under development will also offer high speed – but it's not designed as a LAN, and will likely have an effective range of only a few meters. WPANs are really designed for ad-hoc, temporary use, but they do overlap with WLANs to some degree. We'll return to this point below.

- *Performance (Throughput* and *Range)* - While much additional functionality is still being added to the medium-access control (MAC) layer of 802.11, significant evolution in the physical (PHY) layer has drawn equal if not more attention. The acceleration of basic rates from the original one- and two Mbps of the 1997 standard to 1999's 11 Mbps (802.11b) and 54 Mbps (in 802.11a and 2003's 802.11g) dramatically accelerated end-user demand. While 802.11g is an interesting "mid-life kicker" for .11b, the push is on to exceed 100 Mbps, and this is the goal of the 802.11n Task Group. And we see no practical upper bound on the demands for throughput and ever greater range. It should also be noted here that while real-time communications of the voice-over-IP (VoIP) variety require relatively little raw throughput, the time-boundedness of these communications requires greater headroom in available bandwidth so as not to incur unacceptable delays.

- *Price* – The WLAN market is at present driven to a large degree by price, with moderate commoditization at least in the eyes of buyers. This, however, is about to change. While the price/performance ratio is important, the general decline in the price of high-tech products driven by VLSI always returns attention to performance.

While we see no real change in the importance of 802.11 over the next few years, we thus do see a greater emphasis on performance, initially at the expense of higher product cost. While system builders are today somewhat hesitant to add additional functionality of any form if this improvement has a negative effect on cost, we do not believe that this state of affairs will last much longer. Performance has traditionally been used as a differentiator in the world of wired LANs, and great improvements in performance have usually just preceded rapid market expansion – consider, for example, the dramatic increase in LAN connections that accompanied the introduction of 100Base-T. It is also important to keep in mind that early 100Base-T NICs were much more expensive than their 10Base-T counterparts, but that prices declined quickly with successful products also maintaining backwards compatibility with the earlier 10Base-T functionality. Some vendors have used higher performance, even at a higher price, as a differentiator to gain dominant market share – a good example being 3Com's legendary 3C509 Ethernet board, which offered much higher performance at a significantly higher price than many competitors, and became a market leader in the process.

Regardless, cost in modern electronic devices is always mitigated, to a great degree, by progress in device integration via VLSI – and MIMO implementations can most certainly take advantage of this effect. So we do not expect any cost disadvantage to last for very long, and the winning of major benchmark exercises to be a powerful lure to at least investigate MIMO.

MIMO can today offer dramatic improvements in throughput over competing WLAN technologies (see Figures 2 and 3).

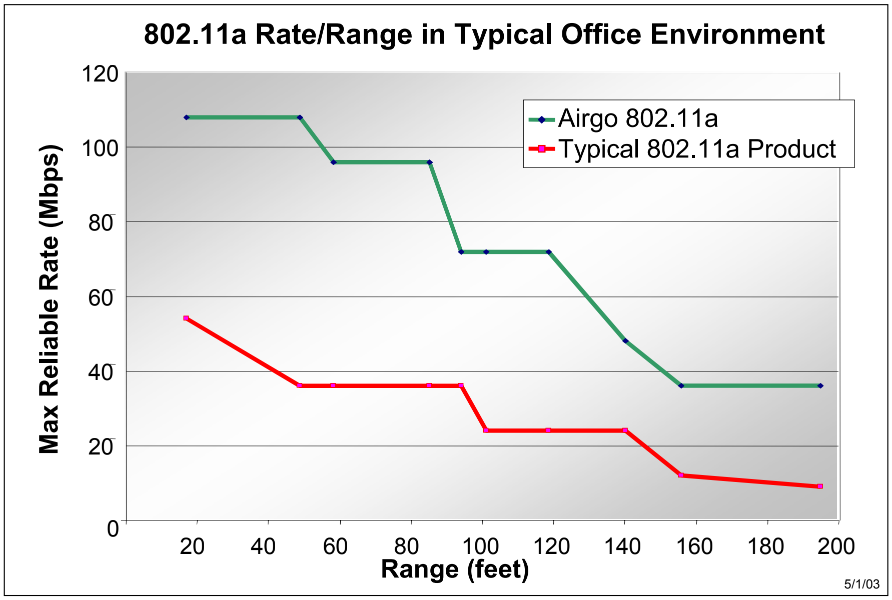

Figure 2: Results of a study done by Airgo Networks, showing throughput and distance improvements using their implementation of a MIMO-based wireless LAN. *Source:* Airgo Networks

MIMO Markets and Applications

In theory, MIMO could have far-reaching impacts in essentially any wireless application, fixed or mobile. A very simple form of MIMO OFDM was successfully deployed (see sidebar, “A Brief History of MIMO”) in fixed microwave line of sight links in the late 70s. With the recent introduction of commercially-viable mass market MIMO WLAN products, we can even expect to see MIMO technology being applied in the future in cellular phone handsets, especially since these devices will also likely contain a wireless-LAN connection for both voice and data use particularly in public-access “hot spot” connections.

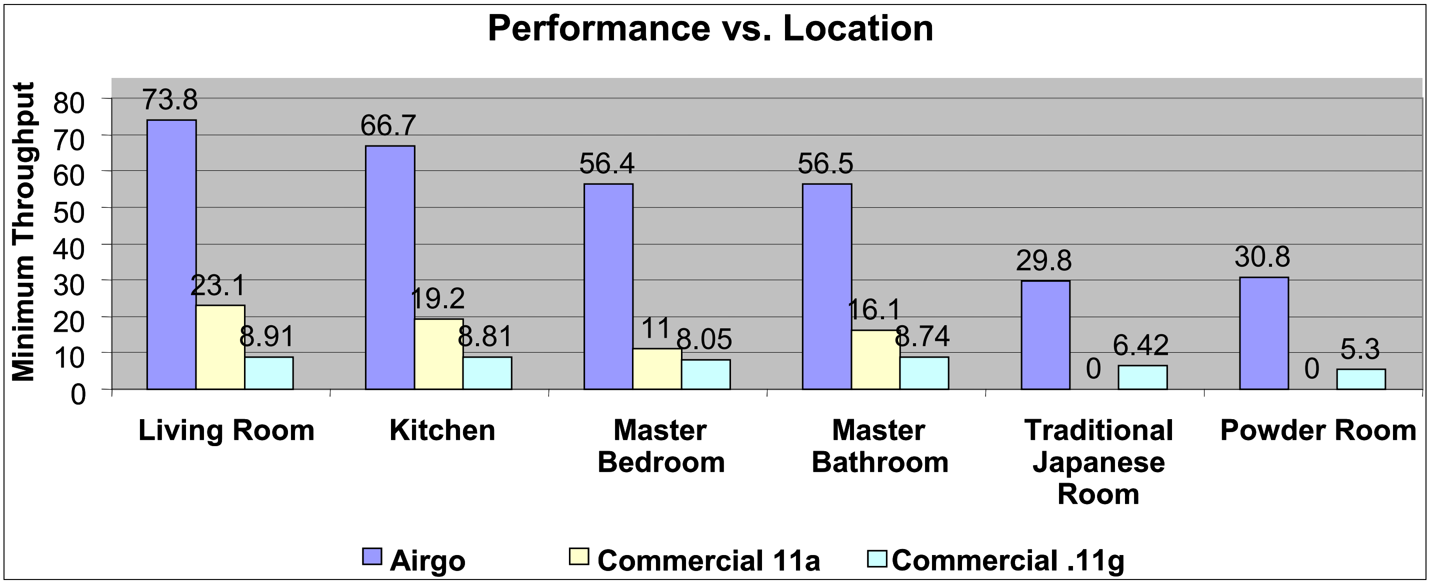

Figure 3: Results of a study of throughput in a Japanese residential application, published by ZDNet Japan at http://www.zdnet.co.jp/broadband/0307/18/lp01.html (in Japanese). *Source:* ZDNet Japan

As we noted above, we expect the demand for effective data throughput to increase over time in wireless systems in a fashion identical to that which has become typical of wired networks. A key benefit of MIMO is in obtaining a greater percentage of total available bandwidth a greater percentage of the time. This will have the effect of evening out the responsiveness of the radio channel, providing a more consistent and wire-like experience. Moreover, as the demand for time-bounded (isochronous) capacity, required for voice and real-time video and multimedia communications, grows with the evolution of applications, MIMO should play an important role in providing a satisfactory experience (as well as greater overall capacity). In summary, we see no market-segment restrictions ahead for MIMO, and we expect it to become a mainstream wireless technology in enterprise, service provider, and consumer-electronics applications. While many expect the ultrawideband approach likely to be commercialized in products based on the upcoming IEEE 802.15.3a standard to dominate in the consumer space, MIMO-based WLANs are likely to be major players here as well.

A final, and important, point: while Farpoint Group has in recent years (with the rapid declines in WLAN capital-goods costs) recommended favoring capacity over coverage, MIMO does shift this strategy somewhat. Since MIMO allows higher capacity and also extends range, it may be feasible in many deployments to begin with a fairly sparse deployment of access points – a strategy reminiscent of the early days of WLANs, when the high price of APs dictated a minimalist approach to infrastructure. MIMO, however, should allow such an approach to be implemented with no degradation of capacity, even as a smaller number of APs covers a larger area that would be possible with other WLAN technologies. And since MIMO-based APs are managed like any others, adding additional capacity as required is straightforward.

Trends and Future Developments

Farpoint Group expects MIMO to play an important *if not critical* role in future wireless LAN systems. As we noted above, the addition of spatial processing is perhaps the best (if not the only practical) way to provide higher performance in terms of both throughput and range. We fully expect the 802.11n standard, which is now under development, to include MIMO processing in its specification. The goal of 802.11n is to define a technical standard for a single WLAN radio channel with performance of at least 108 Mbps (or roughly double that of 802.11a and 802.11g). We also expect that MIMO-based WLAN products with performance of 144 – 200 Mbps will become available in advance of the publication of 802.11n, which is expected to take roughly two years. We expect future MIMO antennas to be built into mobile PCs and PDAs in much the same fashion as today – with the antennas located, for example, in the display portion of the notebook’s housing. Note that a potentially large number of antennas could be placed here, keeping in mind that MIMO performance improves significantly with the addition of more receive antennas.

In early 2000, Farpoint Group selected MIMO as one of the most important technologies for this decade (the others being wideband CDMA, orthogonal frequency-division multiplexing, self-organizing mesh networks, ultra-wideband communications, and software-defined radio). Given the large number of constraints present especially in mobile radio systems, there is always significant debate as to what technology is most appropriate in a given case. We see, however, no reason that MIMO will not grow in influence and importance as radio systems evolve during this decade. With the introduction of MIMO-based WLANs, significantly ahead of our expectations in 2000 as to timeframe, we might add, MIMO technology is certain to become available in a broad range of end-user products. So-called “4G” wireless networks, which we define as all-IP metropolitan- and wide-area networks with support for time-bounded communications, may take advantage of MIMO as well – and we may even see a technological merger of wide-are wireless and wireless LANs at some point. So, needless to say, MIMO remains on our most-important list today - and, as MIMO becomes available in low-cost systems, a fundamental barrier to higher performance has been broken.

Farpoint Group

7 Whippoorwill Lane
Ashland MA 01721
508-881-6467
www.farpointgroup.com
info@farpointgroup.com

The information and analysis contained in this document are based upon publicly-available information sources and are believed to be correct as of the date of publication. Farpoint Group assumes no liability for any inaccuracies which may be present herein. Revisions to this document may be issued, without notice, from time to time.

Copyright 2003 — All rights reserved

Permission to reproduce and distribute this document is granted provided this copyright notice is included and no modifications are made to the original.