WhitePaper

Wireless Without Compromise: Delivering the promise of IEEE 802.11n

Consistent Application delivery for the All Wireless Enterprise.

Executive Summary

The Wi-Fi technology has been improving for several years, to the extent that many students, healthcare workers, and enterprise users now rely on wireless as their primary data connection to the network.

Wi-Fi infrastructure is already a $1 billion market with annual growth in double digits, but most of this growth has come from a few key sectors: manufacturing, retail organizations, hotels, universities and schools. Adoption to date in enterprise offices has been slower, as many CIOs and end users still regard a wireless connection as inferior to a wired Ethernet connection. This is changing with IEEE 802.11n, whose additional capacity has the potential to displace wired networks. When deployed using appropriate technology, 802.11n can enable an all-wireless workplace.

Few network managers will choose to rip out existing Ethernet wires and move to Wi-Fi, of course. However, the all-wireless option is already competitive with wires in greenfield sites that lack Ethernet infrastructure, and 802.11n will rapidly lead to Wi-Fi becoming the dominant mode of access even for users who have the option of wired Ethernet.

Rather than replace existing wires, 802.11n will avoid the need for new edge Ethernet cabling and edge switches that would otherwise be needed to support additional users. The transition to wireless will happen gradually, coinciding with upgrades of PCs and Ethernet edge switches. There are already more laptops shipped than desktops. All enterprise laptops will soon include 802.11n, leading to widespread adoption over the three-year PC upgrade cycle. Similarly, Ethernet switches usually last around five years. And while mobility needs to be supported, network budgets haven't necessarily increased to match. As the next Ethernet switch upgrade cycle comes around, more network managers will consider using the available budget for deploying wireless and extending the useful life of existing wired capacity.

With the deployment of security protocols backed by the Wi-Fi Alliance, it is already accepted that a well-designed Wi-Fi network is as secure as a wired LAN connection, if not more so. Likewise, millions of Wi-Fi phones are in use worldwide, demonstrating the maturity of multimedia over Wi-Fi technology. The latest Wi-Fi advance, 802.11n has now proven that wireless can offer higher performance than most wired Ethernet connections: 802.11n access points currently support data rates up to 300 Mbps, superior to most common 100 Mbps Ethernet connections.

This 5x increase in speed over legacy Wi-Fi standards eliminates the last serious hurdle to adoption of the all-wireless workplace, where no cables need be run to individual desks and workstations. As a result of 802.11n, wireless will soon become the edge of the enterprise network.

The primary benefit of 802.11n is its superior performance, which also means superior range. When deployed with Meru's award winning Fourth Generation wireless architecture[1] using Channel Layering, 802.11n enables connections at much higher data rates with pervasive coverage and reliability exceeding that of Ethernet. Meru has already supported all-wireless office efforts in deployments worldwide. Combining Meru's proven architecture with the increased speed and range capabilities of 802.11n, it is possible for wireless LANs to become the primary network for most enterprises. Meru's ability to provide the same performance in a single channel as traditional WLAN systems provide using three or more means that organizations can effectively deploy 802.11n in the 2.4 GHz band, something that no other vendor can do without suffering significant performance degradation.

Meru's spectral efficiency has another big advantage: Multi-radio APs can use multiple channels for additional capacity or redundancy, something not possible with other systems. This is particularly important at 2.4 GHz., but it is also necessary to gain the full benefits of the increased spectrum now available at 5 GHz. Unlike competitors using legacy technology, Meru makes it possible for enterprises to pervasively deploy secure wireless voice and data networks as the primary connection for mission critical activities.

802.11n Overview

The IEEE 802.11n standard is a huge step forward for the wireless LAN industry. 802.11n updates nearly every aspect of the technology to offer dramatically improved throughput, range and coverage. Before 802.11n, Meru Networks had already distinguished itself with an innovative architecture for deploying high density wireless LANs that deliver high performance and application stability using the IEEE 802.11a/b/g standards. It shipped the first enterprise 802.11n access points in 2007, and will continue to guarantee application delivery within the 802.11n standards framework.

IEEE 802.11n is a set of specifications for high throughput enhancements to the 802.11 wireless LAN standard. It includes a number of improvements at both the Physical (PHY) and Media Access Control (MAC) layers. Ultimately, when all of the 802.11n enhancements are used in an ideal network supporting only 802.11n stations, Physical layer data rates of 600 Mbps per radio will be possible. This is more than a tenfold improvement over 802.11a/g, but the real performance increase can be even greater as the MAC layer efficiencies mean that more of this bandwidth is available to real applications.

Status of IEEE 802.11n and Wi-Fi Alliance

Wireless LAN silicon vendors are already shipping chipsets that support 802.11n. These have quickly become a standard feature on enterprise laptops. Following Meru's lead, many other infrastructure vendors have released 802.11n-based APs, while the chipsets are also finding their way into other devices such as home routers and Wi-Fi phones. Despite this widespread industry support, the IEEE 802.11n specification is not officially a standard yet; formal ratification of 802.11n is expected in June of 2009 .[2]

However the technical details of the 802.11n standard have been finalized for some time and formal IEEE standardization is not a critical requirement for compatibility. In the early days of 802.11 networks, many products that claimed to be compatible with 802.11b did not work together, thanks to the large number of options in the standard. To address this, the industry created the Wi-Fi Alliance (http://www.wi-fi.org), a global, non-profit association of more than 300 companies devoted to promoting the growth of wireless Local Area Networks (WLANs). The Wi-Fi Alliance's testing and certification programs help ensure the interoperability of WLAN products based on various IEEE 802.11 specifications.

The Wi-Fi Alliance's tests are always based on IEEE 802.11 standards, but it doesn't always wait for an official standard. In June 2007, it began a certification program based on Draft 2.0 of the IEEE 802.11n standard. "Draft n" products are certified to be interoperable based on a subset of the 802.11n specification, removing the last barrier to 802.11n adoption in the enterprise. 802.11n Draft 2.0 is already becoming the de-facto standard for high-performance WLAN equipment, with the Wi-Fi logo guaranteeing interoperability.

Customers should always look for the Wi-Fi certification logo before purchasing wireless LAN products to ensure interoperability. Meru's AP300 family of 802.11n Access Points are Wi-Fi CERTIFIED™ for 802.11n draft 2.0 [3]. Its recently announced AP440 will be certified prior to general availability. .However, certification does not guarantee performance, so users also need to investigate the particular features of different vendors' products and the deployment architectures these enable.

In developing IEEE 802.11n Draft 2.0, the basis for the Wi-Fi Alliance 'Draft-n' certification, the IEEE made optional a number of important features of the original 802.11n specification. The full standard will likely include options for up to 4 transmitting and receiving antennas, as well as support for more spatial streams. However, Wi-Fi Alliance 'Draft-n' certification covers 3 antenna chains and 2 spatial streams. When combined with the new 40 MHz channel size and MAC improvements, this delivers up to 300 Mbps raw data rate. When configured properly, the usable aggregate throughput is around 180 Mbps per radio.

The 802.11n standard is also designed to be backward-compatible with 802.11a/b/g: legacy clients can connect to a an 802.11n network, while 802.11n clients can connect to legacy networks. However, legacy connectivity is always at legacy speeds, and some vendors don't support it at all, necessitating a forklift upgrade. 802.11n is also the first standard to offer both 2.4 GHz and 5 GHz. options, meaning that some networks or clients may not be backward-compatible with all three previous standards.

IEEE 802.11n Features and Benefits

MIMO

The most significant change in 802.11n is the addition of MIMO to the physical layer. MIMO, Multiple Input Multiple Output, is the use of more than one transmitter and more than one receiver on the same wireless device to increase performance. MIMO introduces a smart antenna system to the standard that uses multiple transmit and multiple receive antennas to improve RF signal quality and increase the raw data rate.

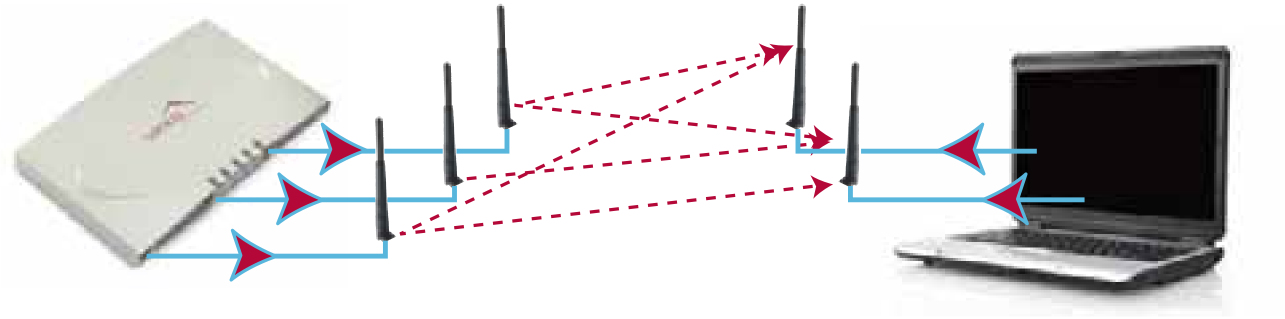

Figure 1: Multiple paths between an access point with three antennas and a client with two

802.11n products are typically described in terms of their MIMO (Multiple In, Multiple Out) attributes, denoted by TxR:S where T is the number of transmit radio chains, R is the number of receive radio chains and S is the number of spatial streams. Most of the 802.11n enterprise APs are either 2x2:2 or 3x3:2 systems while most of the initial 802.11n clients are 2x2:2 systems. Other combinations are also possible. At least two antennas are required for two spatial streams, but many systems have more antennas than the number of spatial streams.

The diversity of transmit and receive radio chains is one way that 802.11n provides improved range and better coverage, though buyers need to keep in mind that the number of antennas is not the only factor determining this. For example, tests show that a Meru 2x2 AP offers better range and coverage than another vendors' 3x3 AP, thanks to Meru's more advanced antenna design. Performance depends on antenna quality, not just the number of antennas.

802.11n also implements another MIMO technique called SDM - spatial division multiplexing. Using the multiple transmit and receive radio chains, it is possible to transmit multiple streams of data simultaneously on the same channel, thereby increasing the data rate and overall throughput. Draft n products must support at least two spatial streams. The full 802.11n specification offers options for up to four spatial streams, though as yet no systems with this feature have been produced.

Rather than the number of antennas, the number of spatial streams is the key factor in determining data rate. Assuming a clear signal, a two spatial stream link will achieve twice the throughput of a single spatial stream in the same channel. Each spatial stream provides up to 150 Mbps of data rate, so a system with two spatial streams will provide up to 300 Mbps data rates.

Improved OFDM

The 802.11n PHY increases the bit rate of the channel through improved OFDM. Originally introduced to Wi-Fi in 802.11a and 802.11g, OFDM is Orthogonal Frequency Division Multiplexing, which breaks the data stream up into several sub carriers that are sent in parallel. This allows more data to be reliably transmitted within the same channel size. The highest bit rate for 802.11a and 11g was 54 Mbps, whereas 802.11n achieves up to 65 Mbps in the same 20 MHz channel with a single stream. As with previous 802.11n variants, the highest data rates depend on a clear signal. When signal quality is poor, the system will drop down to simpler modulation with lower data rates.

The 802.11n physical layer also increases data rate by supporting an optional Short Guard Interval which increases the low level symbol rate by an additional 10%. The Guard Interval is a gap in data transmission used to protect a signal against echoes from previous transmissions, so like the improved OFDM it depends on a high signal-to-noise ratio.

Channel Bonding

802.11n allows two adjacent 20 MHz channels to be combined to form a single 40 MHz channel. This effectively doubles the raw data rate possible. A two stream system will achieve 150 Mbps raw data rate in a 20 MHz channel and up to 300 Mbps in a 40 MHz channel. However, increasing the channel size decreases the total number of channels available. There is only one available 40 MHz channel in the 2.4 GHz band (in North America) and up to eleven 40 MHz channels available in the 5 GHz band.

Meru is the only vendor in the industry that support full 300Mbps with 40MHz channels in the 2.4GHz band.

MAC Protocol Improvements

The 802.11 Medium Access Control (MAC) protocol has also been enhanced compared to earlier versions of 802.11. The improvements here do not actually increase the raw data rate, but they do make more of it available to real applications.

Packet aggregation and Block ACK allow more efficient use of the higher data rate PHY. An ACK (acknowledgement) is no longer required for each data frame: The 802.11n MAC allows multiple frames to be transmitted in sequence, without waiting for an ACK before transmitting the next. A single ACK instead covers a block of many transmitted frames. This reduces the protocol overhead and improves performance, particularly for streaming traffic such as video.

There are other protocol optimizations that reduce overhead, though these require that networks consist entirely of 802.11n stations. . For example, Reduced Interframe Space (RIFS) defines a shorter time interval between successive frames on the airwaves, similar to the Short Guard Interval at the lower level.

Overall Performance Impact

he performance improvement of each of the 802.11n mechanisms is multiplied to produce the effective data rate of the system. A single stream of 802.11n in a 20 MHz channel provides 65 Mbps. Two spatial streams doubles this to 130 Mbps. Using 40 MHz channels increases the raw data rate to 270 Mbps, and adding Short Guard Interval results in 300 Mbps.

As with other standards, the 300 Mbps date rate only refers to the physical layer. Some of this throughput is used by higher-level protocols, meaning that applications will see only a fraction of it. However, 802.11n also introduces a more efficient MAC layer, delivering a higher percentage of the raw data rate as useful throughput. An 802.11g system operating at 54 Mbps raw data rate delivers approximately 22 Mbps of throughput. 802.11n on the other hand is able to achieve around 180 Mbps of throughput on a 300 Mbps raw data rate system.

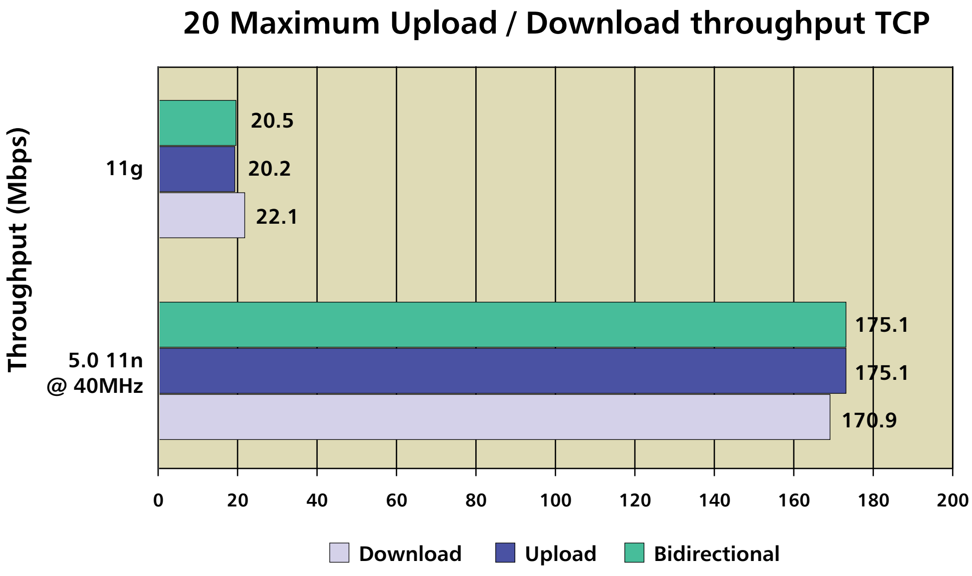

Figure 2: TCP throughput of 802.11g and 802.11n networks compared

802.11n Benefits for Enterprise

Higher Data Rates

Current enterprise 802.11n products that are Draft n certified are capable of 300 Mbps raw data rate over a 40 MHz bonded channel with two spatial streams. The effective throughput for an enterprise class 802.11n client on a properly designed network is about 180 Mbps.

Greater Range and Robust Coverage

Besides increased data rate and throughput, 802.11n improves range and coverage compared to legacy 802.11g and 802.11a. Communication between 802.11n access points and clients is robust due to the adaptive multiple antenna systems used at both ends of the wireless link.

Higher Capacity

Properly designed 802.11n networks can deliver multiple gigabits of useful capacity in a given area. In addition to the increases in raw data rate, the highest capacities of all are achieved using Channel Layering, an architecture made possible by Meru's Air Traffic Control technology.

Compatibility with Legacy Systems

802.11n systems are fully backward compatible with 802.11a/b/g, meaning both that 802.11n clients can connect to legacy 802.11 a/b/g networks and that 802.11n infrastructure can support a mix of 802.11n clients and 802.11a/b/g clients. However this compatibility with legacy systems has a price. Careful planning is required to maximize the performance potential of newer 802.11n systems while also allowing legacy clients to operate on the same network infrastructure.

The 802.11n standard defines three different modes of operation that describe the backwards compatibility of 802.11n networks. These three modes are:

- **High Throughput, Greenfield Mode.** In this mode, only 802.11n clients are supported. The entire network operates at full speed, taking advantage of all the improvements at both the physical and MAC layers. Most users will want to ensure that their 802.11n networks operate in Greenfield Mode as much as possible, though the popularity of previous Wi-Fi standards mean that few users will initially be able to use it on every channel all the time.
- **High Throughput, Mixed Mode.** In mixed mode, the network uses 802.11n when communicating with 802.11n nodes and 802.11a/b/g when communicating with 802.11a/b/g nodes. In this scenario, the 802.11n stations are somewhat faster than the 802.11 a/b/g stations, but they are not able to operate at full 802.11n speed due to signaling overhead. Some Mixed Mode networks can also be slowed by the increased time that legacy clients take to send and receive data, as only one client can access a channel at once. To prevent this, Meru networks use Air-time fairness, a technology that gives all clients equal access to the network unless QoS priorities demand otherwise.
- **Legacy Mode.** This means that the 802.11n equipment essentially emulates 802.11/a/b/g, so data rates are limited to 54 Mbps. However, there is often some improvement over a true legacy network, as 802.11n's multiple antennas can achieve better range and coverage. Because the data rate of

802.11a/g is highly range-dependent – the full 54 Mbps is only available relatively close to the access point, so better coverage can also increase speed and capacity.

Where multiple channels are available, the optimum performance will usually be achieved by confining Mixed Mode to just one channel so that all the others can use Greenfield Mode. However, this is not always possible, as 802.11a may need to be supported in addition to 802.11g, requiring two channels. To ensure that Greenfield Mode is available in both bands, a total of four radios are necessary.

802.11n Enterprise Design Considerations

HIEEE 802.11n has the potential to replace Ethernet as the primary enterprise network access method, thanks to its combination of higher throughput, superior capacity and better coverage than previous 802.11 generations. However, a successful enterprise deployment entails careful consideration of several factors that did not affect previous Wi-Fi technologies. In particular, the architecture of networks using non-Meru technology means that the wider channel size and increased range can cause as many problems as they solve.

AP Coverage is Different with 802.11n

The range of 802.11n networks is much better than earlier standards, but the way that range is achieved is very different and more sensitive to the physical environment. The figure below shows the actual coverage of an 802.11n AP in a particular office environment. Higher data rates are indicated by the darker color. Some areas of high data rate coverage are very far away from the AP and not even contiguous.

In 802.11 a/b/g networks, AP coverage maps are relatively simple. They look like a series of concentric circles emanating from each of the APs, with higher data rate circles closer to the AP. In 802.11n, APs provide coverage that is much less predictable due to its reliance on multipath effects for better signals. Instead of neat, easily arranged concentric circles, coverage cells are spiky and sometimes non-contiguous, leading to more coverage holes and possibly higher co-channel interference.

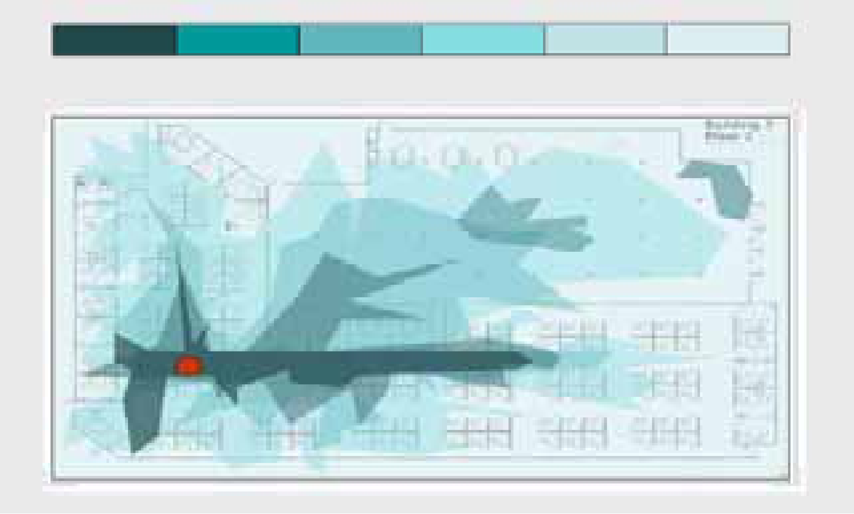

Figure 3: Bandwidth density (shades of blue) surrounding an 802.11n access point (red)

Coverage Planning is Complicated and Unpredictable

Coverage patterns for multiple AP deployments of 802.11n are very difficult to predict because the physical layer technology leverages multi-path: reflections of the RF signal that cause the receiver to see multiple copies of the signal. Multi-path effects were ignored by previous 802.11 systems, but are used by 802.11n to increase both data rate and range. This makes 802.11n systems highly sensitive to their physical environment and means that each deployment is unique.

Vendors using legacy third-generation technology require a microcell architecture, in which adjacent access points are each tuned to a different frequency to avoid interference. Deciding which APs will use which channels in a large network is very complex, and usually left up to automated RF tools that attempt to predict coverage.

Because of the unpredictable, spiky coverage pattern of 802.11n APs, the estimates made by automated tools will be even less accurate for 802.11n than for 802.11a/b/g. As yet, there few 802.11n site survey tools on the market, and they are likely to lead to a greater number of coverage holes than their 802.11a/b/g equivalents. APs will also overlap their coverage with neighboring APs more than desired, increasing interference that degrades performance.

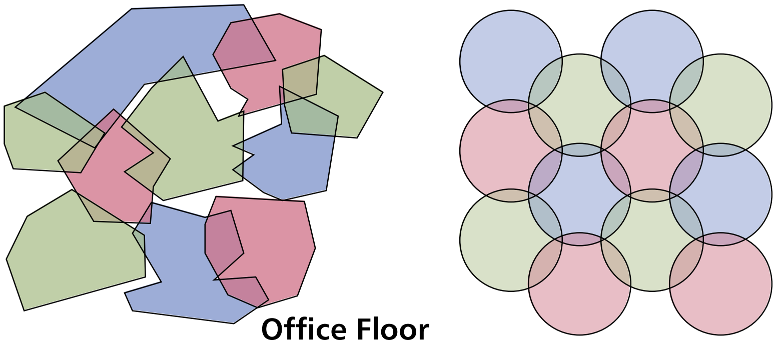

Figure 4: Microcell complexity, colors representing channels. Left: 802.11n; Right: 802.11a/b/g

As the figure above shows, simple design guidelines used for 802.11a/b/g networks such as the common "20% overlap of adjacent cells" estimate are no longer practical with 802.11n due to the irregular coverage patterns. While these result in patterns of microcells such as that on the right for 802.11a/g, the one on the left is more likely with 802.11n, especially in an office environment where closely-packed walls, doors or cube dividers provide multiple radio paths. Predictive coverage planning tools developed for 802.11a/b/g systems do not deliver useful estimates for 802.11n systems, as their calculations are based on RF signal attenuation and cannot account for multi-path at all.

Backward compatibility between 802.11n and legacy 802.11 a/g clients

The 802.11n standard requires that hardware offer backwards compatibility with legacy 802.11 equipment, though users can choose to disable this feature. When enabled, 802.11n devices will operate with degraded throughput mode in the presence of legacy 802.11 a/b/g devices, but still perceive some capacity and range improvements compared to their legacy counterparts.

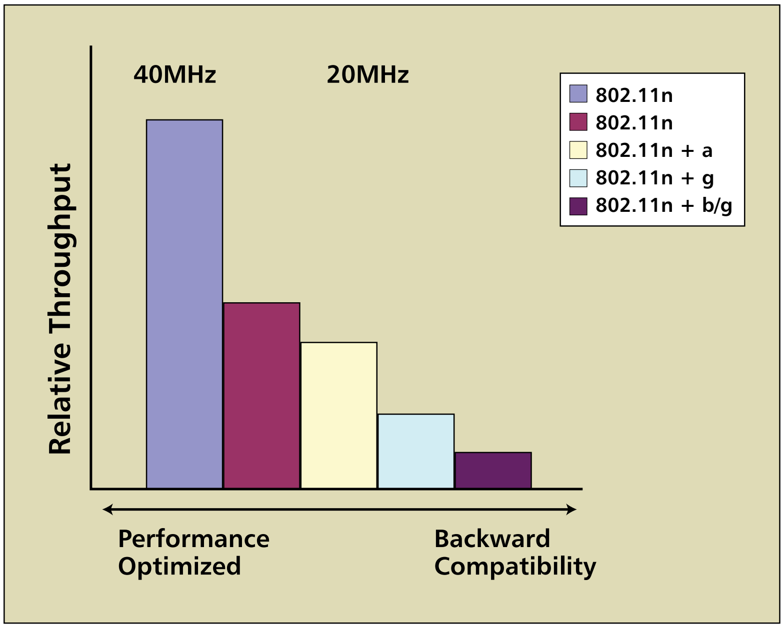

Figure 5: Relative Impact of backward compatibility mode on 802.11n performance

Backward compatibility is a key requirement for any new networking standard, but it has historically reduced the effective capacity of the newer technology. 802.11n is no different from previous standards efforts in this respect. Ideally, to benefit from the full capabilities of 802.11n, the legacy wireless devices should be quarantined to a different channel from 802.11n devices, enabling the .11n clients to operate with no performance impact. This channel isolation is not always possible, as it requires multiple channels and multiple radios. Microcell 802.11g networks already use the entire 2.4 GHz. band to provide just one access layer, leaving no non-overlapping channels free for legacy 802.11b clients. The problem will be even worse with 802.11n, thanks to its wider channel size. In a microcell network, all 2.4 GHz. 802.11n clients may be slowed down by a single 802.11b/g device.

Co-Channel Interference

Radio transmissions can cause interference to other devices at a range beyond their useful communication range. This co-channel interference affects performance, reducing both coverage and data rate of access points based on a microcell architecture. Because the communication range of 802.11n systems is increased compared to previous versions of 802.11, the interference range is also greater. Together with the unpredictable coverage area, this will make co-channel interference an even more serious problem with microcell networks that attempt to upgrade to 802.11n.

Deployment of 802.11n APs

802.11n can be deployed in both the highly popular 2.4 GHz and the 5 GHz spectrum. Though most new clients may eventually use the 5 GHz. band thanks to its much greater number of channels, most .11n clients today support both 2.4 GHz. and 5 GHz. but start scanning at 2.4 GHz. first. If they find a 2.4 GHz. network, they will connect to it, even if it is 802.11b/g only and does not support the higher data rates. This makes support for full-rate 802.11n at 2.4 GHz. critical for optimum performance of new 802.11n clients..

Unfortunately, the microcell architecture used by most vendors places a serious barrier to 802.11n in the 2.4 GHz spectrum, as the highest data rates require the wider 40 MHz channels. Only one 40MHz channel is available at 2.4 GHz., which means microcell networks cannot reliably be deployed at full 300 Mbps speeds in the 2.4 GHz band. will need to leave large gaps between cells to avoid co-channel interference. Without channel bonding, the maximum data rate for 802.11n clients is reduced to 130 Mbps.

Both legacy 802.11 b/g clients and newer 802.11n clients will be on the same channels, further reducing the overall speed and throughput of the network.

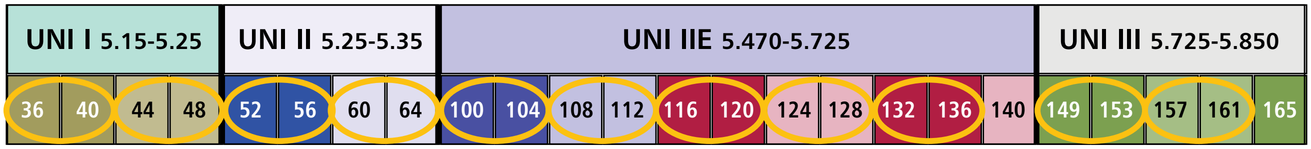

UNI I 5.15-5.25				UNI II 5.25-5.35				UNI IIE 5.470-5.725											UNI III 5.725-5.850				
36	40	44	48	52	56	60	64	100	104	108	112	116	120	124	128	132	136	140	149	153	157	161	165

Figure 6: Total Available 40MHz channels at 5 GHz = 11

The 5 GHz spectrum has eleven non-overlapping 40 MHz. channels available, so a microcells architecture is possible here. However, as shown below it results in higher design complexity for the system, as well as incompatibility with the dominant 2.4 GHz channel plan and network design. The two bands require different channel designs and hence different AP placements.

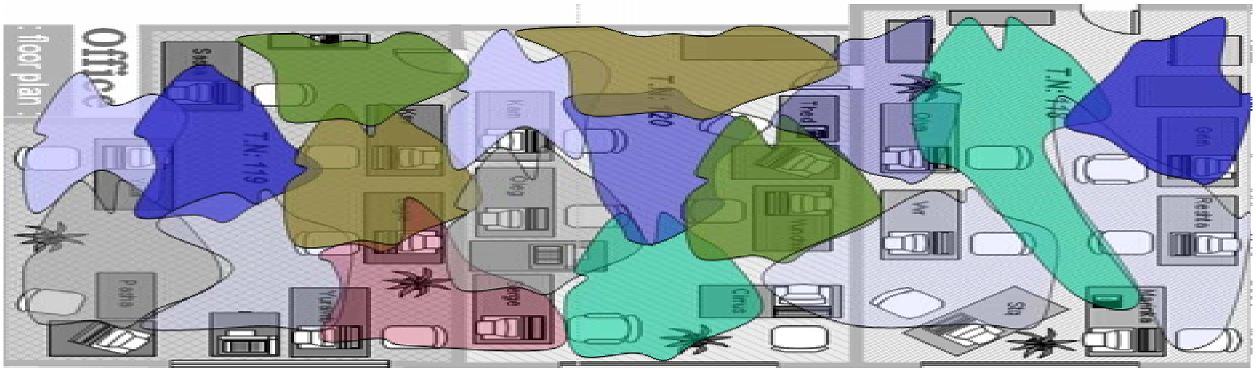

Figure 7: Microcell channel pattern for 802.11n at 5 GHz.

Diversity of Wireless Client Capabilities with 802.11n

The diversity of wireless devices and drivers already causes issues with existing wireless LANs. At 2.4 GHz., slower 802.11b clients throttle the performance of faster 802.11g clients. At both frequencies, “sticky” clients fail to execute a handoff, instead remaining connected to one AP even when they move so far away from it that their data rate drops. Poorly implemented clients or drivers can accentuate these problems, putting a strain on even the most sophisticated wireless networks.

All of these issues are set to get worse with 802.11n. Because the standard is so complex, the difference between the best and basic wireless client and driver implementations can be dramatic. 802.11n offers myriad options in terms of spectrum usage (20 MHz or 40 MHz channels), transmit and receive streams (2x2:1, 2x2:2, 3x2:2, 3x3:2 configurations), encoding (QPSK to 65 QAM, spatial OFDM), and MAC protocol optimizations (TXOP, frame aggregation, block ACK, delayed ACK). A wireless phone or other handheld client may have a 2x2:1 radio operating at low power. Notebook PCs may have a 3x3:2 radio with excellent antennas and much higher power. The perceived range and coverage will be different depending on the type of client even when all support 802.11n.

A single legacy client can have a huge performance impact for all users. As the chart below shows, the difference in the air time required to send the same length packet using different 802.11 standards is dramatic. If all clients are allowed to send the same number of packets, an 802.11n network could spend nearly all its time listening to slow 802.11b transmissions. The worst client will dominate the airwaves and the performance of the entire network will suffer due to the slower clients.

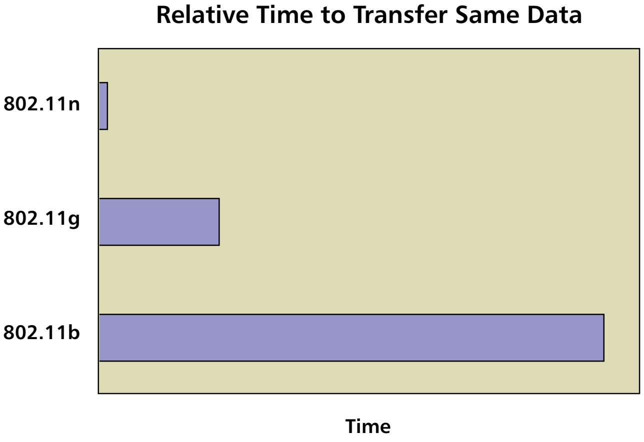

Figure 8: Impact of client diversity on performance

Power Over Ethernet

802.11n access points incorporate multiple radios each with multiple spatial streams, requiring much more processing power than legacy access points. Most LAN edge switches now provide Power Over Ethernet (PoE) as per the IEEE 802.3af standard. DC power is injected at the switch or at a separate inline device, carried over Ethernet cabling to the access point. The nominal limit for 802.3af (at the device) is 12.95W, too low for many 802.11n APs.

Access points can also accept local DC power provided by a plug-in power supply, but many enterprises do not like to use such supplies because they require an AC outlet. This increases the cost and complexity of the installation for each AP. It also means that remote power management is not possible.

A new PoE standard is under development, tentatively called 802.3at. It aims to deliver at least 30W of power to each client device. This standard will not be complete until 2009, and implementation will involve forklift upgrades to LAN edge switches or mid-span injectors. Meru's 802.11n APs will not require this, as they can use the existing 802.3af standard.

Backend Network/Controller support for pervasive 802.11n deployment

With legacy 802.11 a/b/g APs, the throughput per AP is sufficiently low that the only network bottleneck is in the air. For example, a typical network with predominantly 802.11g clients will have a peak effective throughput of 25 Mbps per AP. Assuming an over-subscription factor of five, 200 APs can be handled with a single Gbps Ethernet link. The order-of-magnitude increase in over-the-air throughput also means that the backend network controller must also be able to support a high throughput..

Meru's Fourth Generation WLAN Architecture: Designed for 802.11n

Meru has been shipping its award winning fourth generation WLAN architecture since 2003. A combination of advanced, unique technologies, the Meru architecture fully complies with all Wi-Fi standards including 802.11n Draft 2. When used with standard 802.11a/b/g/n clients, it enables high density wireless LAN deployments in the enterprise with application performance, reliability and security comparable to that of Ethernet.

Fourth Generation: Air Traffic Control

UThird-generation enterprise 802.11 systems use a wireless LAN controller to manage and secure large scale deployments with many APs. Sometimes called switches because early models had to be connected directly to APs (taking the place of Ethernet switches), controllers centralize management, security and IP mobility. The widespread use of controller-based systems has made enterprise wireless LAN deployments practical for enterprise IT departments, but not yet enabled seamless mobility at the high data rates required by multimedia applications.

The fourth generation takes the controller a step further with Meru's Air Traffic Control architecture. In addition to centralized management, security and mobility, the Meru controller governs access to the airwaves in real time. The Meru controller coordinates the activities of every AP and client to reduce interference and ensure efficient operation of the entire wireless network.

In a Meru system with Air Traffic Control, the infrastructure is in charge. Associations between clients and APs are controlled by the network, not by each individual client, resulting in improved mobility and handoffs between APs that are instantaneous and seamless. The Meru controller has global awareness of the entire system and therefore knows which AP provides the best service for each client. Load balancing is automatic.

Air Traffic Control delivers airtime fairness, density control and deterministic QoS. Airtime fairness ensures that each wireless client gets its fair share of the airtime so that high performance clients do not have to wait for legacy clients. Density control ensures that the airtime is not overly consumed with collisions and retransmissions, something that other networks let occur due to differing client implementations of the 802.11 protocols. Deterministic QoS builds on the traffic prioritization of the 802.11e standard, explicitly coordinating voice and data traffic within the network so that they do not interfere with each other.

Virtual Cell Architecture

The conventional approach for deploying multiple access points in an enterprise, shown below, is a microcell architecture that arranges the APs such that adjacent access points are using different channels. The APs then operate independently, with the controller's role limited largely to processing of traffic itself.

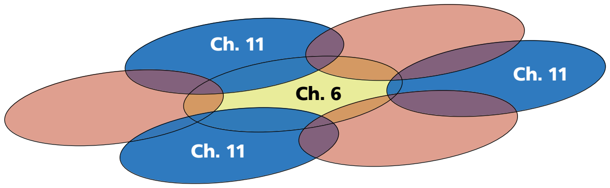

Figure 9: Multiple microcells on different channels

In a microcell system, the power level and channel of each access point must be adjusted to ensure that there is enough overlap to avoid coverage holes. However, the power can not be set too high because this would cause co-channel interference with other access points

In a Meru system, all of the APs use the same channel. The controller coordinates airtime access of clients to ensure efficient operation of the entire network. Each client's traffic is automatically routed through the appropriate access point. Although all use the same channel, multiple APs and clients can transmit without risking interference. Instead of standalone wireless Ethernet hubs, Meru APs operate as a coordinated system of antennas, maximizing parallelism in transmission while minimizing co-channel interference. The entire enterprise can be covered with a single pervasive and homogenous Virtual Cell.

A Meru system is much simpler to deploy because all access points are on the same channel and they can be set at the maximum permitted power. Overlapping coverage areas are an advantage rather than a problem in the Meru system. The Meru controller eliminates co-channel interference by coordinating APs using Air Traffic Control technology.

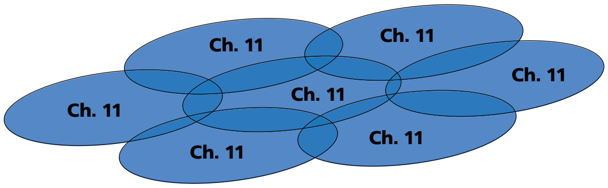

Figure 10: Meru access points all on channel 11

Instead of multiple APs, clients see only one: a single Virtual Cell. This simplifies the network, as the client no longer has to concern itself with decisions about which AP it should link to. Instead, all such decisions are taken by the network itself, with associations automatically load balanced to ensure optimum performance and battery life. Handover is almost instantaneous, as from the client perspective handovers do not occur: The client remains connected to the same virtual AP even as it moves through a large network.

Because APs are all using the same channel, a client can even be connected through more than one AP at once. This makes handoffs much smoother, as a new connection is established before an old one is dropped. It also improves accuracy of location tracking and further extends range. Behavior is more consistent, something particularly important with mobile 802.11n clients as these introduce new roaming algorithms. With Air Traffic Control, client-initiated roaming is unnecessary..

Adding More Capacity: Channel Layering

To increase capacity in any system, more access points must be added. In a microcell system, this means that the access points need to repeat the alternating pattern but with smaller cells. The power level for all APs must be further reduced to prevent co-channel interference.

The Meru system can layer additional Virtual Cells in the same area. Unlike microcells, multiple layered Virtual Cells coexist in the same physical space, each using a different channel. The capacity of the network can thus increase linearly with every new AP or radio added, with the only limit being the number of channels available.

The additional channels can be implemented through either additional APs or additional radios on the same AP. Meru offers omnidirectional APs with up to four radios, an engineering feat still not achieved by competitors.

Network Design Flexibility

Channel Layering effectively multiplies the number of wireless networks by however many radios are installed, though unified management means that the multiple networks can be treated as one. It can seem similar to microcells in that more than one channel is used. The difference is that Channel Layering makes all channels available throughout the network, whereas microcells use a patchwork that makes all but one or two channels unavailable in any particular cell. It has several applications:

- **High bandwidth density**
  A standard AP acts like an Ethernet hub, sharing its bandwidth among all users connected to it. Channel layering allows APs to act as true wireless switches, multiplying network capacity by adding more radios.

- **Redundancy**
  Multiple channels mean that a network is better able to withstand interference. This is particularly important in the narrow 2.4 GHz band, where Wi-Fi networks compete with cordless phones, Bluetooth devices and microwave ovens, but newer phones and radar systems affect 5 GHz networks too. By using an AP with multiple 802.11n radios, such as the Meru AP300 series or AP440, users can ensure that a backup channel is always available when one is blocked by interference.

  Legacy WLAN systems are not capable of this sort of RF redundancy, which is only made possible by using Meru's layered channel approach For even greater redundancy, Meru's AP 440 supports two Gigabit Ethernet uplinks for connections to separate switches and power sources.

- **Reserved bandwidth**
  Voice and data can coexist on a single-channel network, using QoS mechanisms to ensure that voice traffic is prioritized. But for extra reassurance, channel-layering can segregate applications at the physical layer, giving voice and data separate channels. Similarly, channel layering also guarantee bandwidth to particular applications or users.

- **Bigger, faster, more reliable mesh**
  Without access to multiple channels, nodes in a mesh have to buffer traffic until a radio channel is free – a process that can add latency, reducing the number of hops available. Channel layering lets mesh networks separate client connections from mesh itself at the physical layer, making larger meshes possible without sacrificing performance.

  Meru networks can scale up to very dense systems with many Virtual Cells, many access points and many clients. Microcell-based systems cannot, as they already consume multiple channels in an attempt to avoid interference. The Air Traffic Control architecture ensures that the network remains stable and that applications behave as expected even under heavy loads. With Meru's fourth generation architecture, wireless LANs become a credible alternative to wired networks. The all-wireless office can be a reality.

Meru's Unique Architectural Benefits for 802.11n

Meru's Virtual Cell architecture offers a number of advantages compared to conventional micro-cell enterprise deployments and addresses many of the deployment issues that arise with 802.11n.

Simpler Network Planning

Deploying Meru's Virtual Cells greatly simplifies 802.11n network planning. The Meru APs can be deployed on the same channel at high power.

Because of MIMO, the shapes of 802.11n cells are much less predictable than those of 802.11a/b/g cells, making coverage holes and interference more likely in a microcell network. This is not a problem with Meru's Virtual Cell technology, in which the coverage areas of all APs merge to form a seamless blanker. Coverage holes are less likely; if they do occur, adding a new AP will plug the hole without requiring the network to be reconfigured.

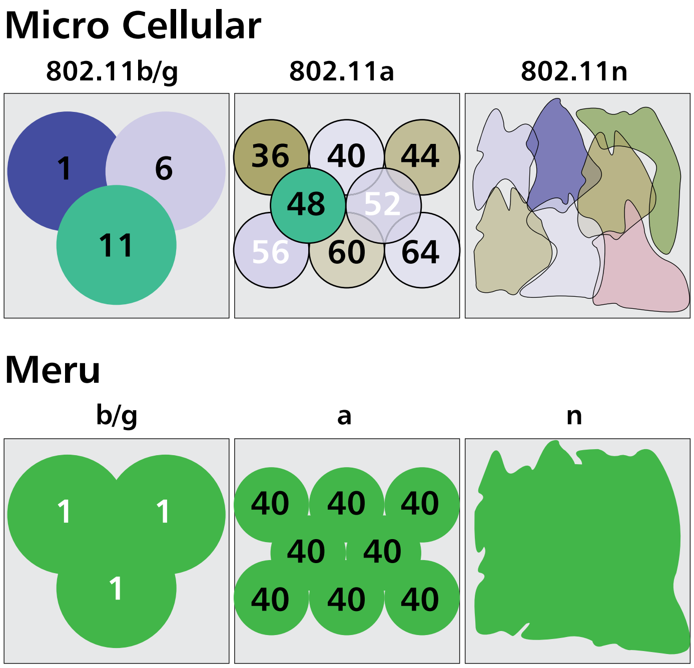

Figure 11: Multiple microcells (top) contrasted with Meru's single-channel Virtual Cell (bottom)

Meru turns overlapping AP coverage areas from a problem into a solution. Instead of interfering with each other, the neighboring APs help each other, providing additional coverage in areas where the signals from one alone would be weaker. The centralized, real-time Air Traffic Control system replaces expensive site surveys and planning.

Air Traffic Control also enables APs to broadcast at maximum power, improving range and extending coverage into areas that microcells would leave as holes. If coverage gaps are found, they can be plugged simply by adding more APs without additional planning or reconfiguration.

Client Investment Protection

Meru's single channel approach also helps with the challenge of migrating from legacy 802.11b/g clients. This is because the 2.4 GHz band has only 3 non-overlapping 20 MHz channels. In a conventional micro cellular deployment as shown below, all of these channels would be consumed providing coverage throughout an office. When upgrading to 802.11n, the entire system must run in mixed mode or suffer from co-channel interference. New 802.11n clients that use 2.4 GHz will have to share the network with legacy clients, operating at lower performance than their full potential.

In a Meru system, legacy clients and 802.11n clients can be separated on different channels that cover the same area. The 802.11n Virtual Cell can use two non-overlapping channels to run in GreenField Mode at 40 Mhz, while a separate, 20 MHz channel serves legacy 802.11b/g clients as shown below.

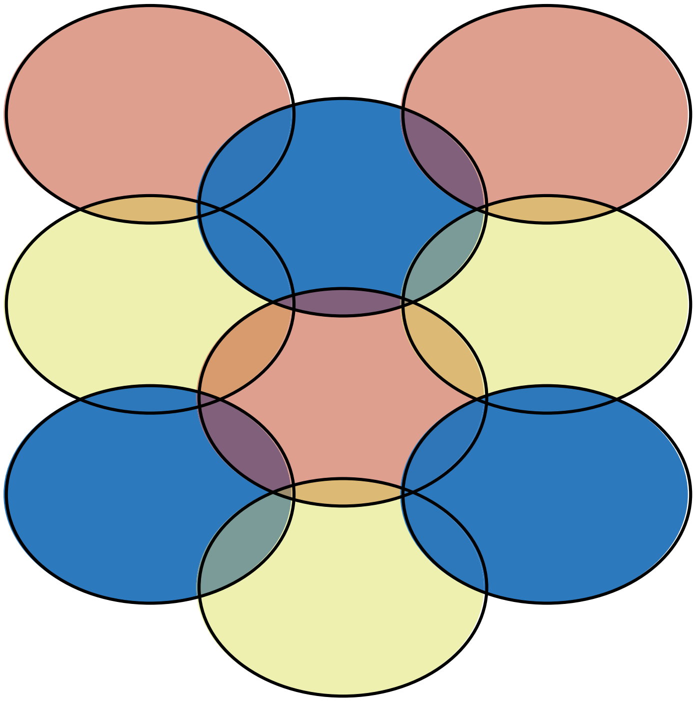

Figure 12: Microcell network consumes all three 2.4 GHz. channels

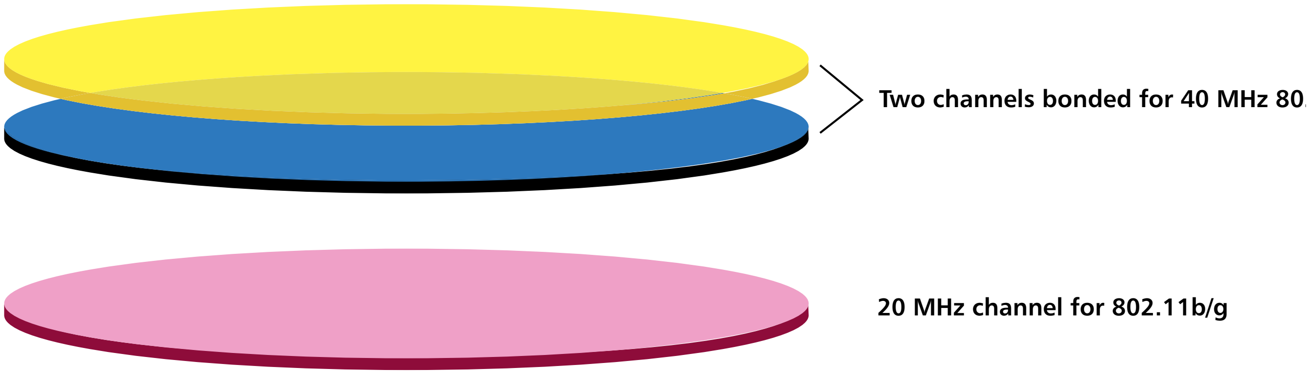

Figure 13: Channel Layering in 2.4GHz

Extending the same concept, it is possible to use a 20 MHz channel at 5 GHz for 802.11a clients and a 40 MHz channel for 802.11n clients operating at 5 GHz.

Delivering Maximum Capacity

Layering Virtual Cells can also be used to add capacity to the network. Each Virtual Cell consumes only one channel, leaving the other channels available for expansion. The alternating channel design of microcell networks consumes all of the available channels to essentially deliver one channel of capacity in each area.

Highest RF Reliability

The extra capacity enabled by Channel Layering also makes a network more reliable, as interference usually affects only one channel at once. Microwave ovens were once the main interference sources, but now many other devices compete for scarce 2.4 GHz spectrum. Wi-Fi networks must contend with DECT phones, Bluetooth devices and neighboring companies' access points. Interference affects 5 GHz networks too, thanks to new cordless phones and to FCC rules that give priority for many 5GHz channels to radar systems. The interference is set to get worse with 802.11n, as channel bonding to 40 MHz makes each channel twice as vulnerable.

With Meru's four-radio AP440, users have three other options if one channel is suddenly blocked. For maximum reliability, spare Virtual Cells can be kept on hot standby in both the 2.4 GHz and 5 GHz bands. Air Traffic Control automatically switches clients over to another Virtual Cell if one becomes unavailable.

Meru's 802.11n Product Line

Meru offers a comprehensive product line that is 802.11 Draft 2.0 compliant and Wi-Fi Alliance Draft n Certified.

AP 300 Family

The Meru AP300 family of dual-radio 802.11 a/b/g/n access points represents the next generation of wireless edge devices, providing superior, reliable support for high-capacity data and toll-quality voice. All can run two 802.11n radios at full speed within the power limits of the current 802.3af standard, meaning that no new injectors or DC power cabling are necessary.

AP 302

Configuration: Two a/b/g radios. Each radio is dual band (2.4GHz and 5.0GHz) and upgradeable to 802.11n.
Benefits: Serve legacy a/b/g clients on both radios today; upgrade to 802.11n in future.

AP 310

Configuration: Single a/b/g/n radio Dual band. Serves all clients either on 2.4 or 5 GHz 802.11n.
Benefits: Lower cost 802.11n solution, ideal for layering in .11n on a legacy a/b/g network.

AP 311

Configuration: Two radios (1 a/b/g/n + 1 a/b/g) Each radio dual band (2.4 GHz and 5 GHz).
Benefits: Serve legacy clients on a/b/g radio and 802.11n clients at 5.0 GHz or 2.4 GHz on second radio. Ideal for migrating from legacy a/b/g to 802.111n at lower cost. Upgrade second radio to 11n in future as needed. Because the AP's radios have built in support for 11n, the AP311 also provides rogue detection for 802.11n devices, something that other a/b/g APs cannot.

AP 320

Configuration: Two radios (both 802.11a/b/g/n radios.) Each is dual band (2.4G Hz and 5.0 GHz).
Benefits: Serve 802.11n clients on 5.0 GHz. or 2.4 GHz. on both radios. Ideal for deploying layered 802.11n channels to achieve maximum capacity and throughput.

Meru AP440

The Meru AP440 is the first AP to include four 802.11n radios. Each supports the full 40 MHz channel size, allowing a total capacity of up to 1.2Gbps. This makes it an effective way to replace Ethernet, combining the flexibility of wireless with the performance and stability of wires. It is the only way to take maximum advantage of the new standard's high performance without compromising security, reliability or interoperability.

The AP440 features two 5 GHz radios, one 2.4 GHz radio and one dual-band radio. This makes it flexible enough to handle multiple deployment scenarios. Channel Layering can be used to increase capacity, to provide backup Virtual Cells, or to segregate legacy clients or particular applications. Alternatively, the dual-band radio can be used for security monitoring, scanning both bands for potential attackers.

The AP440 integrates three advanced wireless technologies into a single compact unit with no obtrusive external antennas. It also features USB connectivity for devices such as sensors or spectrum analyzers

Four radios is too many for the current 802.3af standard, so the AP440 can be powered in three different ways: with a separate DC power supply, using the new 802.3at technology, or using two 802.3af connections in parallel as it has two Gigabit PoE ports.

Meru MC 5000

The Meru MC5000 Chassis is a modular, telco-grade, multi-gigabit, five-blade chassis controller for up to 2000 APs (per chassis.), with 4 Gbps throughout scalable to 20 Gbps. The MC5000 was designed to efficiently handle the increased bandwidth requirements of 802.11n. For enterprises with gigabit Ethernet at the network edge and multi-gigabit switching capability at the core, the MC 5000 is a centralized controller that delivers the performance required for large scale 802.11n deployments.

Meru MC 4100

The MC4100 is a self-contained controller that supports up to 4 Gbps of encrypted throughput. Designed to work alone or in concert with other Meru controllers, the MC4100 allows network administrators to securely manage and easily control their wireless networks. It combines seamless mobility through Meru's Virtual Cell technology with the capacity necessary for 802.11n networks.

Meru 802.11n Advantages

The 802.11n specification improves Wi-Fi performance in four main ways. Though all have been broadly adopted across the industry and Wi-Fi Alliance certification gives users a guarantee that every manufacturer's equipment will all work together, not all 802.11n products are equal. The Alliance tests for interoperability, not performance. Meru's unique Channel Layering, Virtual Cell and Air Traffic Control technologies are critical for getting the most out of an 802.11n system.

The AP440 is the only access point on the market that can provide four independent channels of 802.11n, but Channel Layering and Air Traffic Control are also supported on other Meru products. This lets customers mix and match AP440s with previous Meru APs to fit the needs of specific regions within a network. In contrast, most other vendors use older microcell technology whose drawbacks are magnified by 802.11n.

Multiple Input, Multiple Output (MIMO)

Radio signals have been taking multiple routes between transmitter and receiver in all wireless systems since 19th Century telegraphs, something that has traditionally caused problems due to destructive interference between the signals. The effect is most significant in indoor areas, where there are many different ways for radio waves to reflect off, refract through or diffract around obstacles such as walls and cube dividers.

Also known as spatial division multiplexing (SDM) or path diversity, MIMO turns what had been a bug into a feature by using the multiple paths to carry different streams of data. The 802.11n Draft 2 standard says that devices must support two spatial streams, requiring at least two antennas at both the transmitter and the receiver, though additional antennas can be used to improve range. Because both streams are using the same frequency, the antennas must have a very high gain and be able to discriminate accurately between the two signals.

With Meru's Virtual Cell architecture, all access points in a network are already operating at the same frequency. Meru has more than six years' experience in multi-radio antenna design and single-channel RF processing, giving it unrivalled expertise in the type of engineering needed in MIMO systems. This is why Meru was first to ship the first enterprise 802.11n products and why Meru hardware consistently outperforms competitors in independent tests.

The AP440 includes four antennas for spatial diversity, all of them omnidirectional and available for use by all four of the radios. The shared antennas ensure that every radio has maximum spatial diversity, improving the signal-to-noise ratio and minimizing overlap between the two data streams. Omnidirectional antennas mean that all four channels are usable all around the access point, in contrast to competitors' sectored antennas that limit each channel to a specific region.

Channel Bonding

In 802.11a/b/g, all channels are the same 20 MHz width. Upgrading to 802.11n adds a new 40 MHz mode, effectively tying two channels together. The new 40 MHz channel more than doubles the data rate, but it also doubles the likelihood that interference will affect the channel and halves the total number of channels available. This makes radio redundancy and efficient use of spectrum more important than ever.

Meru's Virtual Cell technology already uses spectrum more efficiently than rival architectures, occupying only one channel to provide the same coverage for which microcells consume at least three. This is particularly important at 2.4 GHz, where the total bandwidth is not enough to accommodate three non-overlapping 40 MHz channels. But it also matters at 5GHz, enabling 40 MHz Channel Layering for extra capacity or redundancy. The AP440's four radios mean that users have up to three backup channels available when one is temporarily blocked by interference.

Some 802.11n clients may support only the 20 MHz channel size while others will be capable of the full 40 MHz. With four radios, these different types of clients can be kept separate by dedicating one radio in each band to each channel size. Alternatively, Meru's Air Time Fairness ensures that 40 MHz and 20 MHz clients can coexist on the same channel without all being slowed down.

Modulation Gains

The increase in data rate between 802.11b and 802.11a/g resulted mostly from more efficient modulation and encoding, packing more bits into each radio wavelength. The comparable improvements in 802.11n are not as dramatic, but they still boost performance by about 20%. The drawback is that they depend on a clear signal. Like previous technologies, 802.11n will drop down to lower speeds when a signal is weak or drowned out by interference.

Signals from Meru APs are intrinsically stronger than those from competitors', thanks to their higher antenna gain and transmit power. APs in a microcell network must reduce their power output to avoid interfering with each other, something Meru avoids through its unique Air Traffic Control technology. As well as enabling Virtual Cells, this lets each AP transmit at the maximum power allowed by the FCC, ensuring that the higher data rates are available over a wider area.

When interference occurs, the AP440's radio redundancy and Channel Layering means that clients don't have to slow down to more easily understood encoding. They can simply move to one of the other three available channels.

MAC Layer Efficiencies

Compared to wired LANs, the original 802.11b protocol was very inefficient. Though the physical layer operated at 11 Mbps, TCP/IP applications often saw under 2 Mbps of usable TCP/IP throughput because of protocol headers and error correction. 802.11g improved on that, but more than half of the available capacity was still wasted. With 802.11n, Wi-Fi moves closer to Ethernet in efficiency, offering about 180 Mbps at the application layer. As with Ethernet, some of the remaining overhead is unavoidable due to the TCP/IP stack itself.

Like the improved modulation, these efficiencies depend on a clear signal. Some of the overhead in the 802.11 protocols is really built-in error correction, so eliminating it is only worthwhile when clients can be assured a clear, error-free link. This requires APs that can safely transmit at a high power and that have other channels available when interference occurs.

Higher throughput also puts additional strain on back-end infrastructure. Meru's MC5000 controller supports up to 20 Gbps of encrypted traffic, the highest capacity in the industry. Like other Meru controllers, this can be connected to the AP440 and other APs through any standard Gigabit Ethernet switches.

Conclusions

IEEE 802.11n products, can deliver performance and reliability equivalent to 10/100 wired Ethernet connections, creating a path for the all-wireless workplace: one with no voice data cables to the desk.

Standardization

The IEEE 802.11n Draft2.0 standard and the Wi-Fi Alliance Draft-n certification appear to be stable. The benefits of 802.11n in terms of capacity, performance and reliable coverage have already been well-established by Meru in early campus wide deployments.

Benefits

Global all-wireless-office adopters have seen millions of dollars in increased productivity due to fixed-mobile convergence, as well cost savings from avoiding edge Ethernet upgrades. As many people are adopting a mobile work style the wireless network is following the wireless phone as the primary method of network access.

Risks

The risk of early Draft-n equipment requiring upgrades is very small. The performance of the Draft-2.0 standard will stand through mid 2009, while subsequent products will incorporate more features and offer higher performance with full backwards compatibility.

(Endnotes)

1 Gartner describes the Fourth Generation WLAN architecture as one in which the WLAN controller coordinates air time access across multiple APs in a system, reducing co-channel interference and delivering more consistent performance. A third-generation architecture is one in which clients must decide which AP to connect to and where adjacent APs cause co-channel interference, necessitating a microcell architecture where each AP must be tuned to a different channel from those of nearby APs,

2 The current 802.11n schedule can be found here: http://grouper.ieee.org/groups/802/11/Reports/802.11_Timelines.htm

3 The certifications are listed here: http://certifications.wi-fi.org/wbcs_ViewCertificate.php?product_id=5617

Glossary

802.11n Draft 2

The version of the IEEE 802.11n specification selected by the Wi-Fi Alliance for interoperability certification. Widely adopted across the industry, 802.11n Draft 2 uses MIMO, channel bonding, improved modulation and a more efficient protocol stack to achieve data rates of up to 300 Mbps. Despite the name Draft, the specification is equivalent to a standard, as all equipment with Wi-Fi Alliance certification has been tested for interoperability.

Airtime Fairness

Method of governing access to the airwaves so that all clients are able to transmit for the same amount of time, meaning that performance is higher for 802.11n users than for legacy clients. Without airtime fairness, slower clients can hog the airwaves as they take longer to transmit each packet.

Air Traffic Control

Meru technology that exercises a high degree of control over all transmissions within a wireless network. Unlike superficially similar technologies from other vendors, Air Traffic Control governs is in charge of client connections, not just access points, enables it load balance connections between APs.technology coordinates uplink and downlink transmissions on a single 802.11 channel in such a manner that the effects of co-channel and adjacent channel interference Air Traffic Control can also coordinate AP transmissions so well that co-channel interference is eliminated and all access points on a network can share a single radio channel.

Antenna Diversity

A Technique that improves radio performance by using multiple antennas. Antenna diversity is necessary for MIMO but useful for improving signal quality for any radio type. This means that all other things being equal, 802.11n access points can even improve range or throughput when used purely in legacy 802.11a/b/g mode.

Block ACK

Bundling together of several acknowledgement signals, meaning that a transmitter does not have to wait for an acknowledgement after sending each frame. This is used in 802.11n to reduce protocol overhead and increase the effective data rate.

BSSID (Basic Service Set Identifier)

A 48-bit number used to identify an 802.11 wireless service. In most enterprise Wi-Fi networks, a service means an access point and the BSSID is the same as the AP's MAC address. In a Virtual Cell architecture, all APs appear to have the same BSSID and MAC address and the client sees only one large virtual AP.

Co-channel Interference

Radio interference that occurs when two transmitters use the same frequency without being closely synchronized. Legacy wireless systems cannot achieve this kind of synchronization, so access points or cell towers that transmit on one channel must be spaced far apart. The result is coverage gaps that must be filled in with radios tuned to another channel, resulting in an inefficient and complex microcell architecture. Air Traffic Control technology avoids co-channel interference by tightly synchronizing access point transmissions, enabling that adjacent APs to use the same channel.

Channel Bonding

The combination of two non-overlapping 20 MHz. channels into a single 40 MHz. channel, doubling the amount of data that can be transmitted in a given time but halving the number of available channels. Along with MIMO, it is a key innovation in the 802.11n standard.

Channel Layering

Wireless LAN architecture in which several Virtual Cells are located in the same physical space but on non-overlapping channels, multiplying the available capacity. This additional capacity can be used for redundancy or to support higher data rates or user density. Channel Layering can also segregate different applications or client types, for example keeping 802.11n clients on separate channel from legacy 802.11b/g so that the .11n network can operate at maximum capacity. Channel Layering can be enabled through multiple radios on one AP or by using multiple AP close together, so the total capacity is limited only be the number of non-overlapping channels available.

Channel Reuse

A pattern in which different APs can use the same channel. In microcell networks, such APs need to be placed far apart to avoid co-channel interference, meaning that contiguous coverage requires multiple channels. In networks using Air Traffic Control technology, the same channel can be reused throughout the network, meaning that only one channel is required and others are left free for other purposes.

Controller

Appliance that manages a wireless network and (usually) aggregates traffic. Controllers were introduced with third-generation wireless LAN systems as a way to manage access points. In a fourth-generation system, the controller also governs client transmissions, deciding which AP each client is connected to. Controllers are sometimes referred to as switches because early versions took the place of Ethernet switches and had to be connected directly to access points, though this is now rare as most can now be placed anywhere in a network.

Fourth Generation

Term coined by analyst firm Gartner to describe a wireless LAN system in which the controller governs handoffs, such as one utilizing Virtual Cells. This is contrasted with third generation (“thin access point”micro-cell architecture) systems, in which the controller is only responsible for managing access points and clients must decide for themselves when to initiate a handoff. Second generation (“fat access pointwithout controller”) systems lacked a controller altogether and were designed for standalone operation, whereas the first generation lacked any enterprise management features.

Guard Interval

A gap in transmission during a radio signal to ensure that a receiver is not confused by echoes or other multipath effects. In 802.11a/g, the interval is 800 ns. The 802.11n specification introduces an optional shorter guard interval of 400 ns, which data rates but requires a clear signal.

Meru APs can also be considered “thin”. We call it coodinated AP.

Handoff

The transfer of a link from one access point to another as a client moves through a network. In legacy microcell networks, Wi-Fi clients are responsible for handoff, meaning that the quality of the link and the overall network performance is dependent on each client’s implementation of 802.11 roaming algorithms. In Virtual Cell networks, the network itself governs handoffs as clients remain connected to a single virtual AP.

Microcell

Wireless architecture in which adjacent APs must be tuned to different, non-overlapping channels in an attempt to mitigate co-channel interference. This requires complex channel planning both before the network is built and whenever a change is made, and uses spectrum so inefficiently that some co-channel interference still occurs, especially at 2,4 GHz.. Microcell architectures were common in 2G cell phone systems and legacy wireless LAN systems. They are not used in 3G cellular networks or in wireless LAN systems that use Air Traffic Control, as these allow all access points to share a single channel.

MIMO (Multiple In, Multiple Out)

Technique that increases wireless data rates by sending different data streams over different physical paths at the same frequency. Current systems based on 802.11n Draft 2 support two spatial streams, while the final 802.11n standard is likely to support have an option of supporting up to four. Because the number of paths depends on the presence of obstacles which may be moving, MIMO makes 802.11n coverage very hard to predict.

Modulation

The process of encoding data into radio waves. One of the ways that 802.11n increases data rate is by using more complex modulation techniques, something that requires high signal quality. All 802.11 systems will automatically drop down to simpler modulation (and thus lower data rates) when reception is poor.

Multipath

The multiple routes that a radio signal can take to reach its destination as it reflects off, refracts through or diffracts around obstacles. Multipath has traditionally been viewed as a problem because copies of a signal would interfere with each other, causing effects such as “ghosting” of broadcast TV. In 802.11n, it is turned into a strength by MIMO, which sends different signals via different routes to increase the overall data rate.

OFDM (Orthogonal Frequency Division Multiplexing)

The way that 802.11n, 802.11a and 802.11b subdivide each radio channel into narrower frequency ranges. It achieves higher performance than the spread spectrum techniques of 802.11b by splitting a data stream into multiple narrowband streams that are sent in parallel. Each one can use simpler modulation than would be required for the complete data stream, making the signal less vulnerable to interference or multipath effects.

Protocol Overhead

The amount of a link’s capacity that is used by communications protocols and thus not free for use by application-layer data itself. Legacy 802.11 standards have a high overhead of about 60%, pushing the real throughput of 802.11a/g down to about 22 Mbps. The 802.11n protocol stack has a lower overhead, delivering real throughput of up to 175 Mbps.

Roaming

The process that takes places as a client moves between the coverage areas of different APs, necessitating a handoff. In microcell Wi-Fi networks, roaming can be a complex procedure that risks dropped connections and drags down network performance, as the client is forced to decide when to disconnect from one AP and search for another. In networks using Virtual Cell technology, the infrastructure controls roaming, automatically connecting each client to the optimum AP.

Spectral Efficiency

The ratio of data rate to radio spectrum usage. Of the new technologies introduced in 802.11n, MIMO and new modulation techniques represent higher spectral efficiency, but channel bonding does not (as it uses twice as much spectrum.) With any Wi-Fi variant, a Virtual Cell is much more spectrally efficient than a microcell architecture, as the microcells consume at least three non-overlapping channels to provide the coverage that a Virtual Cell offers with just one.

Single Channel

Term sometimes used to describe a network in which all access points operate on the same channel, such as one using Air Traffic Control technology. Single channel operation is more spectrally efficient than a microcell architecture and necessary for the use of Virtual Cells and network-controlled handoff. However, it also possible to implement single channel without implementing Virtual Cells.

Sticky Client

A client device that remains associated with one access point even as it moves into better range of another, reducing performance and battery life. Sticky clients result from different implementations of roaming algorithms among client or driver vendors and can only occur when clients are responsible for handoff. The Virtual Cell architecture eliminates sticky clients by placing the network in charge of handoffs.

Virtual Cell

Wireless LAN architecture in which a client sees multiple access points as just one, all sharing a single MAC address, BSSID and radio channel . Air Traffic Control. Because clients remain connected to the same virtual AP as they move through a network, no client-initiated handoffs are necessary. Instead, the network itself load balances traffic across APs, maximizes bandwidth, simplifying network management and conserving radio spectrum for redundancy.

VoFI (Voice over Wi-Fi) or VoWLAN (Voice over Wireless LAN)

Voice over IP links that run over a wireless network. VoIP does not usually require high data rates, but it stresses wireless networks in other ways by demanding low latencies and smooth handoffs. In addition, no 802.11n phones yet exist, as most handsets are too small to accommodate MIMO's multiple antennas spaced a wavelength apart. This means that 802.11n networks running VoFI must have a way to deal with 802.11b/g clients.

Wi-Fi

Brand name for wireless LANs based on various 802.11 specifications. All products bearing the Wi-Fi logo have been tested for interoperability by the Wi-Fi Alliance, an industry group composing every major 802.11 client and infrastructure vendors. The Wi-Fi Alliance's tests for 802.11n require adherence to 802.11n Draft 2, as well as portions of related standards such as 802.11i which governs security.