Industry Leading Performance with Cisco's Aironet 3600 Series Access Point

The Cisco® Aironet® 3600 Series Access Point introduces several innovations that result in industry-leading performance. This white paper explains why that matters, and how we do it.

Market and Technology Overview

As three-spatial-stream devices come to market and as the proliferation of iPads and other one- and two-spatial stream devices continues, it is critical to maximize performance for all devices. Not doing so could result in a slower network and in slower application performance for all devices.

Cisco's Aironet 3600 Series Access Points were designed to give best-in-class performance to all devices, allowing optimal network performance and investment protection so enterprises can effectively accommodate any device allowed on the network, regardless of how many spatial streams it has.

The keys to the best-in-class performance of the 3600 Series are:

- The first enterprise-class four transceiver 802.11n multiple-input multiple-output (MIMO) design (4x4:3)
- Cisco ClientLink 2.0 beamforming that works with all 802.11n clients
- Cisco CleanAir technology to manage interference

How does it all work? It's worth starting with a bit of history. The first generation of 802.11n devices that came to market beginning a few years ago were able to support a maximum data rate of 300 Mbps. This data rate could be achieved by running two spatial streams, each carrying 75 Mbps of data per 20 MHz of spectrum, over a double-wide 40-MHz channel. The formula was 75 Mbps x 2 streams x 2 channels = 300 Mbps. Note that in order to support two spatial streams bi-directionally, a minimum of two MIMO transceivers was required at both ends of the link (access point and client).

Very recently, a newer generation of 802.11n devices has emerged in the market, and these devices can support up to three spatial streams. This means that the theoretical maximum data rate that can be achieved is now 75 Mbps x 3 streams x 2 channels = 450 Mbps. In order to achieve this maximum data rate bi-directionally, both ends of the link (client and infrastructure) must support three spatial streams, which in turn requires that both ends must have at least three MIMO transceivers.

Maximum speed is important, but we also need to consider how often it will be achieved. It turns out that getting to the 450 Mbps maximum data rate is not easy and requires careful design, as we'll see later in this white paper. For now, keep in mind that current solutions that use 3x3:3 architecture (three transceivers, three spatial streams) on **both sides** of the link will **rarely** achieve 450 Mbps in real-world applications and will have significant performance degradation beyond a short distance. In real-world scenarios, where clients are not all 10 feet from the access point, the fourth transceiver on one end of the link is required to provide the necessary reliability for three spatial streams to work. And the logical place to put the fourth transceiver is naturally on the access point, where size and power are of less concern than on battery-powered clients.

In the end, the fourth transceiver on the access point gives extra decibels of link margin, which then translates to better performance. As this paper will explain, in the uplink direction (client to access point), the extra receiver allow for MIMO equalization gain. This means that 450 Mbps can be achieved at twice the distance of existing solutions (30 feet instead of 15 feet).

In the downlink direction, the extra transmitter allows for beamforming to the client. The Cisco implementation of beamforming known as ClientLink 2.0 was designed to work with **all 802.11n clients** - both those that support standards based beamforming and those that don't. This includes support for beamforming to clients which support three, two, or even one spatial stream.

So, while a 3x3:3 architecture offers only a nice peak-data story, the 4x4:3 3600 Series offers both high peak-data rate and real-world utility. Of course, building a 4x4:3 solution requires a custom chip-set and additional engineering, so it takes a customer-responsive vendor such as Cisco to deliver it.

For real-world benefit, three-spatial-stream devices must work with usable range for a typical enterprise use. To get this, you really need the Cisco Aironet 3600 Series. Competitive designs deliver data rates of 450 Mbps only at a range of 15 feet or less. And since access points are usually ceiling mounted, a range of 15 feet does not cover 15 feet of floor space. Because the signal must also travel from the ceiling to the typical height of a device, the area covered at the maximum rate by current designs is closer to 12 feet from the access point. With the 3600 Series, the usable range for three spatial streams at 450 Mbps goes to 30 feet. This doubling in range actually results in a 500 percent increase in the coverage area that can achieve 450 Mbps.

The features in the 3600 Series also yield excellent benefits with older 802.11n clients. The increase in the maximum rate coverage area for single-stream clients is Y, and for dual-stream clients, it is Z. These benefits are very important because it will take a long time for all clients to be three spatial stream, and phones will likely stay at one spatial stream.

The Gory Details: Uplink

Although three MIMO receivers are theoretically sufficient to handle three spatial streams, the problem is that it provides no redundancy against channel fading or any of the inevitable hardware impairments: its operation is disappointingly short range or erratic in practice. For this reason, Cisco leapfrogged the products that depend on pure spatial multiplexing, and delivered hybrid spatial multiplexing and diversity with the addition of a fourth receiver.

Cisco's solution is akin to GPS positioning: to determine the latitude and longitude of a GPS receiver, three GPS satellite signals are necessary. More than three GPS satellites are required for superior location accuracy, and even entry-level GPS products can track 12 satellites or more.

The benefits of adding the fourth receiver to support three spatial data streams come from redundancy gain and diversity gain.

Redundancy Gain

The signal captured by the extra antenna provides some redundancy gain. As in linear algebra, the fourth receiver provides four equations while there are only three unknowns. The extra equation provides an additional dimension of freedom when resolving the transmitted signals. One its own, redundancy gain increases range by about 10 percent.

Diversity Gain

Diversity gain is more important. Due the channel fading, one antenna may see a deep fade, resulting in a very low signal level received on that antenna. For three spatial data streams, one needs a minimum of three strong signals so that three data streams can be resolved; if just one of them experiences deep fading, the MIMO detection will fail. By comparison, with 4x4 MIMO, if even one antenna experiences a deep fade, there are still three good signals from which three data streams can be resolved. In actual implementations, the signals from all four antennas are carefully weighted according to quality, ensuring that the signals with better quality are better utilized. 4x4 MIMO has markedly better performance in fading channels due to its balanced hybrid of spatial multiplexing for speed and diversity gain for robustness.

The technical term for the way the fourth receiver gets used is **MIMO equalization**. MIMO equalization is the comprehensive means to make best use of the received signal, whether it be transmit beamformed, space-time block-coded, spatially expanded, or unimproved. For MIMO equalization, having more receive chains helps the most, with the biggest gain coming from one extra receive chain, and diminishing returns thereafter. Accordingly, a 3x3 access point is a reasonable choice for receiving two spatial streams, but not for three spatial streams.

The proof of the benefit can be seen in Figure 1, which shows rate at different ranges in the uplink direction. From the blue curve, you can see that with a 3x3:3 client sending to a 3x3:3 access point, 450 Mbps can only be achieved only at short distances from the access point.

On the other hand, by adding a fourth receive antenna at the access point (that is, the black curve), 450 Mbps becomes usable out to 30 feet, and the third spatial stream adds value all the way out to 40 feet. This behavior is well matched to the coverage range of access points in typical enterprise deployments.

Figure 1. Uplink PHY Data Rate (After Retries) Versus Range

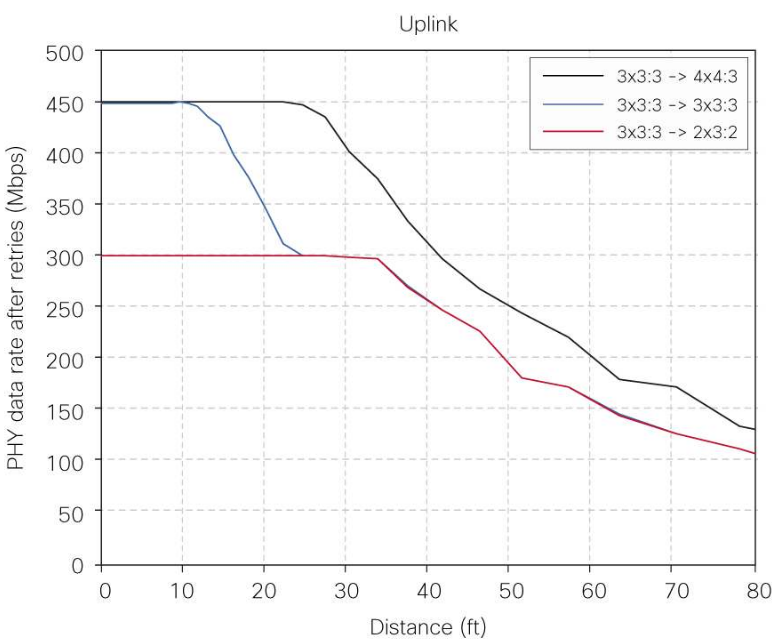

The Gory Details: Downlink

To get a parallel benefit in the downlink direction (from the fourth transmitter), Cisco created ClientLink 2.0, which is a scheme to fully exploit the four MIMO transmit chains. It combines beamforming with spatial multiplexing to improve the speed and reliability of downlink traffic. Even when clients transmit frames with fewer spatial streams than antennas, ClientLink 2.0 utilizes its advanced, patent-pending algorithms to provide full beamforming gain, whether clients help to sound out the channel or not. Furthermore, Cisco access points are fully engineered to avoid any dependence on clients for calibration assistance. However, if such client assistance is available, ClientLink 2.0 is ready to take advantage of it.

Consider the case of a client limited to three transceivers and trying to receive three spatial streams. Detection is very complicated and sensitive to poor signal quality, so any help from the access point is very welcome. As Figure 2 illustrates, using their four transmit chains, Cisco Aironet 3600 Series Access Points have the degrees of freedom to form three spatial beams (where each beam carries one stream of data) and direct them to each of the receive chains at the client. These beamformed signals add up in-phase at the receiver and thereby combat channel fading.

Figure 2. Beamformed signal projection with 4th transceiver

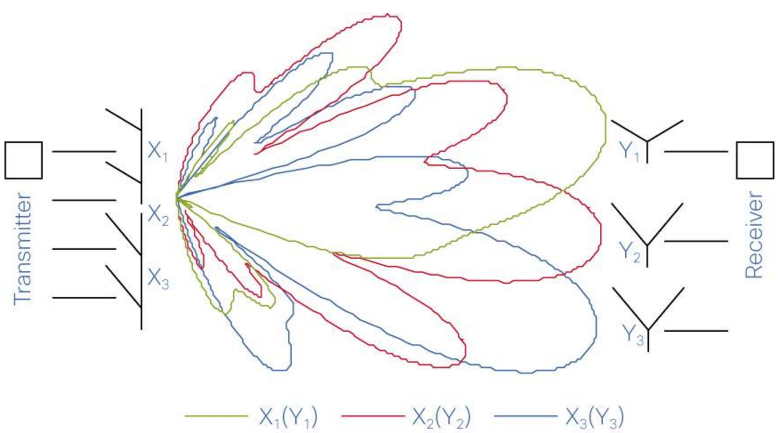

The benefits of ClientLink 2.0 are twofold. First, all four transmit chains are used even when the number of the data streams is less than four. This results in some gain on the signal level due to more transmit power. Second, each data stream is sent from all four transmit chains collectively, resulting in significant diversity gain (one data stream will not be completely wiped out if one or two antennas are in a deep fade).

As Figure 3 shows, the result is evident in the improvement to the achieved data rate at range, which also correlates closely with the reliability of each rate at that range. Assuming a typical enterprise environment, and at a typical enterprise range of 30 feet, a 4x4:3 architecture operates at 420 Mbps whereas 3x3:3 can't do better than 300 Mbps.

Figure 3. Downlink PHY Data Rate (After Retries) Versus Range with Three Receive Chains

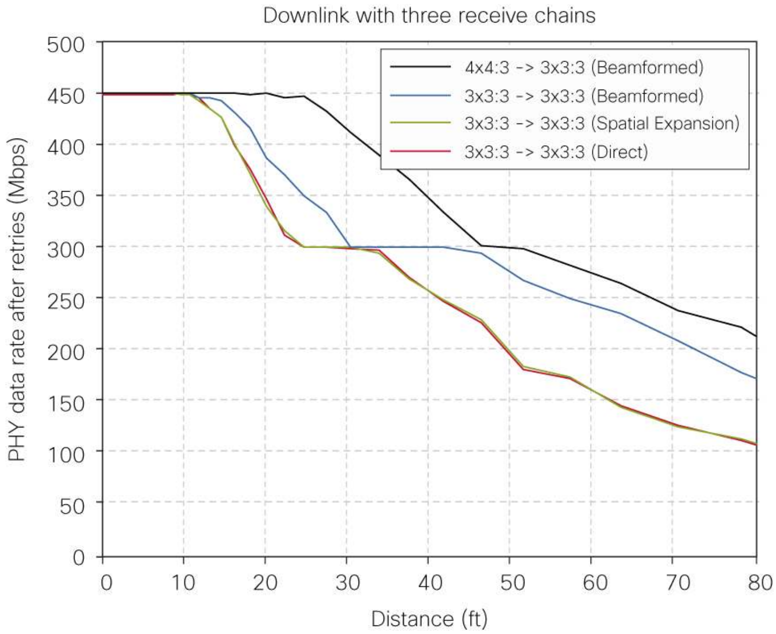

Figure 3 shows the gains from the extra transmit chain. A conventional 3x3:3 access point (“3x3:3 -> 3x3:3 (Direct)” or “3x3:3 -> 3x3:3 (Spatial Expansion)” - the red or green curves respectively - achieves 450 Mbps **on average** only out to 12 feet in line of sight (and only 9 feet horizontally), yet a 4x4:3 access point with transmit beamforming (“4x4:3 -> 3x3:3 (Beamformed)”, the black curve) can deliver 450 Mbps as far as 25 feet (24 feet horizontally). We see that just one extra dimension of redundancy, backed by sophisticated signal processing and RF expertise, is a huge boost. Figure 3 assumes a 20-ms delay between sounding and transmission of the data packet.

Some Important Notes on the Performance Charts

- Since wireless and Wi-Fi devices are variable, Figure 2 and Figure 3 present typical results using a mix of industry-standard assumptions and Cisco experience - 802.11n channel model D: 802.11n impairments (which are somewhat optimistic); 15-dBm transmit power (reasonable for clients and enterprise access points using radio resource management R); 2-dBi antennas at the access point; a 6-dB noise level; and Cisco’s measured path-loss model (which is more accurate than the 802.11n path-loss model).
- Figures 2 and 3 present the **average data rate**. At the specified distance, some clients achieve higher data rates but many others get lower data rates. Also, after MAC overheads, goodput is generally lower, coming in at about 70 percent of the quoted PHY data rates.
- The distance on the x axis is a line-of-sight distance. Thus, for example, when an access point is mounted on the ceiling 8 feet above a client, a distance of 10 feet on the x axis represents a horizontal range of only 6 feet. In other words, 450 Mbps with three receive antennas has very modest utility. Further out, the line of sight and horizontal distances are about the same: for example, 30 feet line of sight is 29 feet horizontally.

Conclusions

A 3x3:3 architecture offers a peak rate of 450 Mbps, but nothing more. To make that peak rate to be achievable over a useful range, the Cisco Aironet 3600 Series provides a higher level of engineering and the custom silicon that only a larger vendor like Cisco can. On the uplink, four receive antennas and an optimized MIMO equalizer extend the range of 450 Mbps out to 25 feet, and sustain rates above 300 Mbps out to 40 feet. Even at longer ranges, the additional diversity provides a significant advantage over a 3x3:3 design. On the downlink, the barrier to success is even higher. An additional transmit chain is necessary, but in addition client-agnostic beamforming, such as ClientLink or ClientLink 2.0, is essential to deliver the higher rates over the important ranges.

Appendix 1: Primer on MIMO and Associated Performance-Enhancing Techniques

The recent amendment 802.11n introduced a set of new features that significantly improved WLAN range, reliability, and throughput. Among the new features, the most profound is the rich multiple-input multiple-output (MIMO) technique that allow multiple data streams to be sent simultaneously (Figure 4).

Figure 4. How MIMO Systems Work

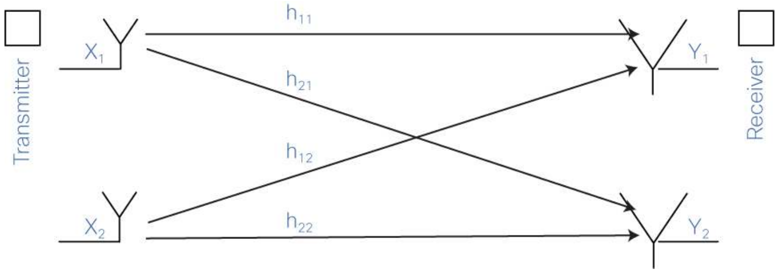

MIMO is a flexible technique that can be used in different ways. At one extreme is pure transmitter and/or receiver diversity, while at the other extreme is pure spatial multiplexing. In the middle is a hybrid mode that offers both diversity and spatial multiplexing.

In the case of pure transmitter/receiver diversity, the same data streams are transmitted and received with multiple antennas. Since the same data is transmitted and received as multiple copies, the possibility of errors occurring in the signal is significantly reduced. Thus, the benefit of the transmitter/receiver diversity is to enhance the link robustness.

In the case of pure spatial multiplexing, multiple data streams are transmitted and received simultaneously over the same bandwidth. The data is different between the data streams. Given the same frequency bandwidth and over the same time period, the data throughput increases **N**-fold, where **N** is the number of the data streams. Thus, the benefit of spatial multiplexing is to increase the data throughput.

The hybrid MIMO mode offers the best of both worlds: additional diversity and spatial multiplexing, to improve both link robustness and data throughput. The hybrid MIMO mode is possible whenever there are more transceivers than data streams. For example, one could transmit and receive two data streams between two 4x4 MIMO devices.

For the transmitter diversity and hybrid modes, getting maximum benefits does require the vendor to go beyond the basics, however. Transmit beamforming is the premier technology, followed by space-time block coding (STBC) and spatial expansion. Let’s look at each of these in turn.

Transmit Beamforming

If multiple copies of the same data are sent from multiple transmit antennas, after going through the wireless channel, the multiple copies of the data have different attenuations and phases at the receive antennas. (Think of phase as like the sign of a number, such as +1 or -1, yet also allowing shades of gray). This is because each copy of the data starts from a different antenna and gets reflected by different walls and furniture as it traverses its path to the receive antennas. If there are two paths arriving at a location with equal attenuation, but one path has the opposite phase of the other, they cancel each other out, which is not good if you're trying to receive the signal at that location!

This headache motivates counter-measures. If the phase for each data stream is known, by precompensating the phase at the transmit antennas, the multiple copies of the data should have the same phase at the receive antennas. Instead of cancelling, they constructively combine with each other. This scheme to precompensate the phases so that multiple copies of the data arrive at the receive antennas with the same phase is called **transmit beamforming**.

Transmit beamforming allows the multiple copies of the data arriving at the receive antennas to have the same phase so they add up with maximum reinforcement. Transmit beamforming significantly improves the signal quality received at the receive antennas. Transmit beamforming doesn't just accept the wireless link it is handed - it also changes and improves it.

Transmit beamforming does require information about the phases (and oftentimes attenuations too) between the transmit chains and receive chain(s) - in other words, channel information. There are three ways to obtain this information; two depend on client hardware support based on the 802.11n amendment, and one can work with any client:

Standards-Based Beamforming

There are two types of standards-based beamforming: explicit beamforming and implicit beamforming.

In explicit beamforming, information about the wireless channel is fed back to the transmitter by the receiver. In order for the receiver to measure the wireless channel, the transmitter first sends a special sounding packet from all transmit antennas. The receiver examines the sounding packet at each receive antenna, extracts the wireless channel information, and sends that information back to the transmitter.

Explicit beamforming is an optional mode defined in 802.11n and requires the support of the client (receiver). The channel sounding protocol incurs some overhead. Still, explicit beamforming provides the most accurate knowledge about the whole channel, from all transmit chains to all receive chains. When available, it should be used judiciously.

Implicit beamforming does not require that a sounding packet be sent. Instead, the channel information is obtained by utilizing the symmetry or reciprocity of the channels that is characteristic of Wi-Fi systems. The transmit and receive chains at the access point share the same set of antennas, so when an access point receives the uplink signal from the client, the access point extracts the channel information from the client's transmit chains to the access point's receive chains. Further, due to channel reciprocity, the same channel information applies equally well to the downlink transmit beamforming.

Standards-based implicit beamforming augments these capabilities to address two issues. First, implicit beamforming is most straightforward when a MIMO client transmits out of all its antennas, which is not always the case. When the client transmits out of fewer antennas, the access point cannot measure the wireless channel fully and cannot maximize beamforming gain. Second, implicit beamforming requires that a device's transmit hardware and receive hardware be well matched. According to implementations, this may be possible during manufacturing, or via internal calibration in the field, or in the worst case may require over-the-air assistance from the client. For these reasons, 802.11n defines optional hardware modes where:

- The client can send additional sounding information from its additional transmit chains even if they weren't being used to send actual data. This mode, when supported, has very low overhead.
- The client can assist with calibration of the access point.

Cisco ClientLink and ClientLink 2.0 Beamforming

802.11a/g clients cannot support standards-based beamforming, either explicit or implicit, and many 802.11n clients do not support standards-based beamforming either. For this reason, it is vital that vendors aiming at a comprehensive solution provide beamforming modes that work for any client. Cisco ClientLink and ClientLink 2.0 beamforming do just that.

ClientLink is for single-antenna clients or clients using one spatial stream. The access point can measure the wireless channel any time the client transmits, even it's just one packet, and then the access point uses that information to maximally reinforce the data send-back to the client. For multiple-antenna clients, vendors must use innovative technology such as ClientLink 2.0.

ClientLink and ClientLink 2.0 beamforming offers high gain, works with every 802.11a/g/n client - even 802.11n clients that do not support standards-based beamforming - and incurs no overhead.

Space-time Block Coding

In this basic mode, the same data is sent over multiple data streams from multiple transmit chains and received on one - and sometimes more - receive chains. Although the data in the multiple data streams is the same, each data stream is cleverly arranged to be different from the others through swapping across antennas or across time, or a sign change. These transformations add excellent time diversity without the challenges of beamforming, yet without the maximum gains either: space-time block coding (STBC) can only make the best out of the wireless channel it is given.

STBC is used to improve the link reliability/robustness in products that cannot do transmit beamforming. STBC requires hardware support in both access point and client.

Spatial Expansion

Spatial expansion is another hybrid mode, where the number of data streams with different data is less than the number of transmit chains. In the usual implementation, some of the data streams employ spatial multiplexing, and some streams employ a diversity scheme (that is, a single data stream sent from multiple transmit antennas). Compared to transmit beamforming, spatial expansion is blind; compared to space-time block coding, spatial expansion is weak. Still, spatial expansion does offer a modest throughput gain as well as modest improvements to link robustness.

MIMO Equalization

On the receiving side, the menu of technologies is very short: it includes only MIMO equalization. MIMO equalization is the comprehensive means to make best use of the received signal, whether it is transmit beamformed, space-time block coded, spatially expanded, or unimproved. For MIMO equalization, having more receive chains helps the most, with the biggest gain coming from one extra receive chain, and diminishing returns thereafter. Accordingly, a 3x3:3 access point is a solid choice for receiving two spatial streams. So too is a 2x3:2 access point, since it provides the same number of receive antennas and so identical gain.

Appendix 2: Downlink Beamforming with Various Conditions

Figure 5. Downlink PHY Data Rate (After Retries) Versus Range with One Receive Chain

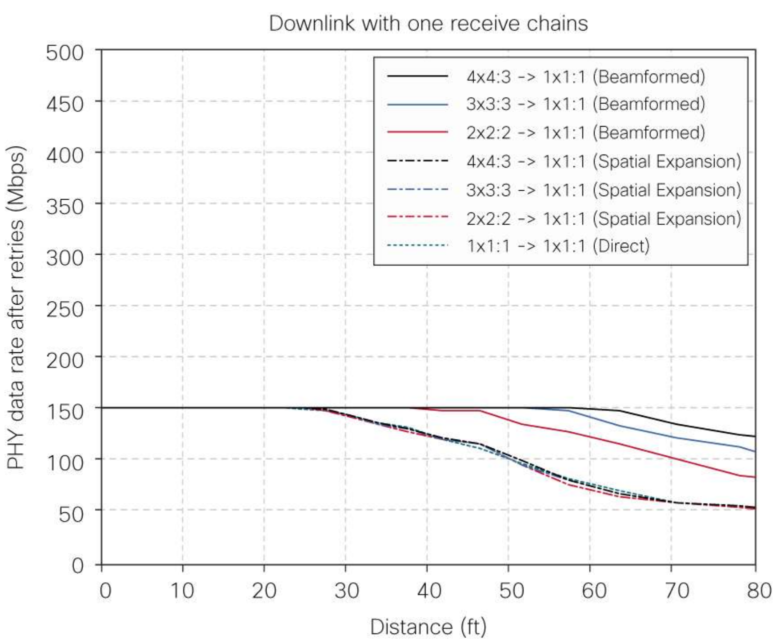

Figure 6. Downlink PHY Data Rate (After Retries) Versus Range with Two Receive Chains

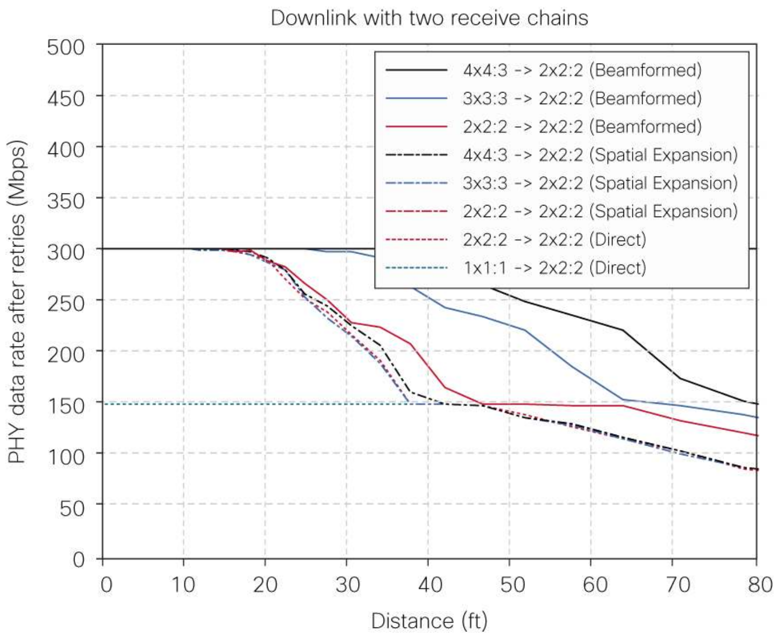

Figure 7. Downlink PHY Data Rate (After Retries) Versus Range with Three Receive Chains

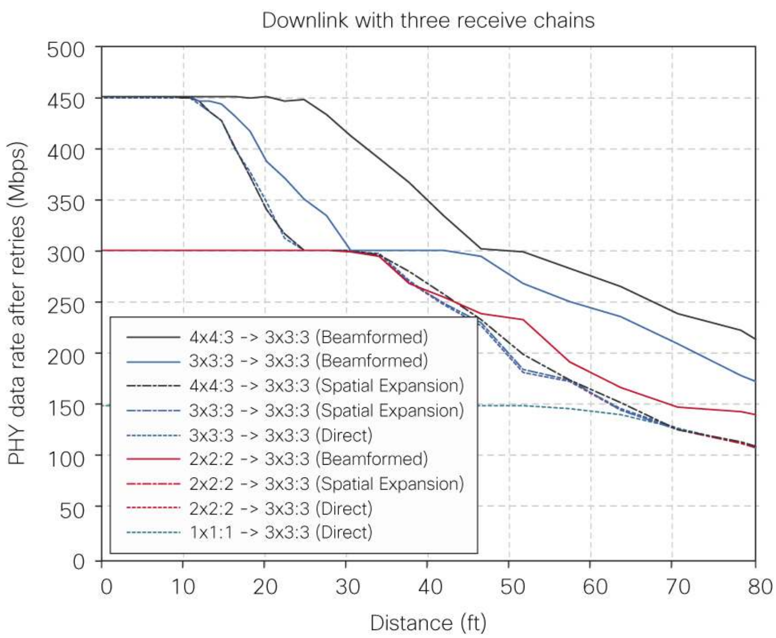

Figure 8. Uplink PHY Data Rate (After Retries) Versus Range

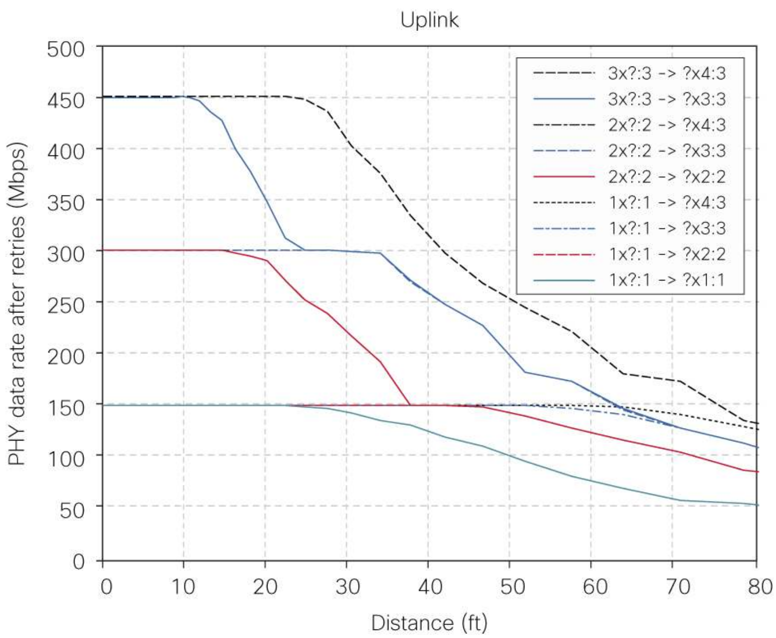

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-688713-00 11/11