bluesocket

The leader in secure mobility® solutions

Demystifying MIMO

A Bluesocket *BluePaper*

Abstract

Multiple Input Multiple Output (MIMO) Technology is poised to make a significant impact in the overall performance of Wireless LANs (WLANs). This paper introduces the concept of MIMO and discusses the architectural differences from traditional Single Input Single Output (SISO) systems. System implementation options are outlined and the advantages and capabilities of both full and partial MIMO implementations are discussed in the context of the architectures. Several examples of the benefits of MIMO and test data are presented. Finally, the implications and timing of IEEE 802.11n are described.

Introduction

MIMO is an acronym for Multiple Input, Multiple Output and is a significant departure in the architecture and technology from the current SISO (Single Input, Single Output) WLAN products. This is accomplished through the implementation of multiple transmitters and receivers in a single WLAN station, AP or client.

A principle problem in wireless systems is multipath interference. This is commonly referred to as reflections, refractions, echoes, coverage holes or “dead spots”. In the cellular world this can best be expressed by having a perfect 5-bar signal and a clear call, then taking a few steps and hearing static. This problem is particularly the case in existing WLANs where systems are characterized, in part, by their tolerance to inter-symbol interference, or delay spread. Delay spread is the difference, in time, between arriving copies of the same signal caused by multipath and is due to path length differences.

Existing WLAN technologies attempt to compensate for multipath by implementing a rake receiver to allow for specific amounts of delay spread. As delay spread increases, the multipath situation worsens and the arriving symbols begin to interfere with each other beyond the capability of the receiver to recover. This results in common complaints to the help desk with the generic “wireless doesn’t work” intermittent reports which cannot be reliably recreated.

Further, the throughput capability of the WLAN with MIMO is extended to be more comparable with wired LANs. The current IEEE 802.11n draft includes link rates up to 600 Mbps. The commonly used value for link overhead in WLAN systems is 40% giving a throughput potential of about 360 Mbps in a single cell. This is an idealized value and is highly dependent on the operating environment and the composition of clients associated with a particular AP.

A basis for MIMO was created at Bell Labs in the 1997 to 2002 period called BLAST for Bell Labs Layered Space-Time. Using MIMO-like techniques, multiple transmitters and receivers, BLAST exploits multipath to gain very high spectral efficiencies (10s of bits/sec/Hz were measured).

The spectral efficiency of a wireless system can be increased significantly through these technologies. However, in a Wireless LAN, the space available for the additional antennas is limited and consequently, three sets has been the chosen number based on the inter-antenna spacing required to reduce co-site interference.

MIMO – The Pros and Cons

The advantages of a MIMO-based WLAN are increased range, throughput and robustness of the data link. In traditional WLAN systems, these objectives are traded off against one another as there is a relationship between data-rate and range. MIMO is not immune to these tradeoffs, but due to the multiple transmitters/receivers and DSP technology, all three of these critical parameters are greatly improved.

MIMO systems demonstrate greater overall throughput using two separate methods. First, the ability to provide a higher data-rate vs. range means improved throughput over a much larger coverage range than traditional SISO antennas. Second, additional throughput can be gained by operating the MIMO AP with 40 MHz channels to increase the available bandwidth for the signal and increase the channel bit rate. While the first technique is applicable to any client device, the second requires a MIMO based client to utilize the additional channels.

The figures below illustrate the overall performance at tested by Tolly Labs of a generic 802.11a/g Access Point, vs. that of an Access Point equipped with MIMO technology:

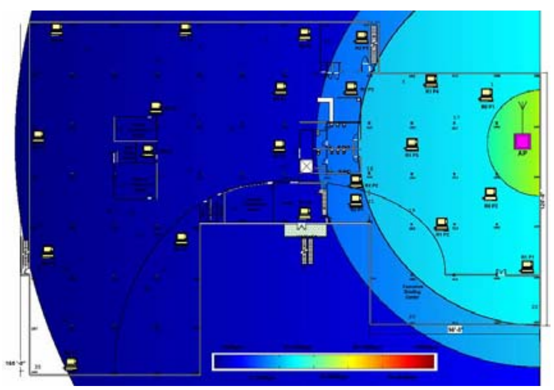

Figure 1 Standard 802.11g AP

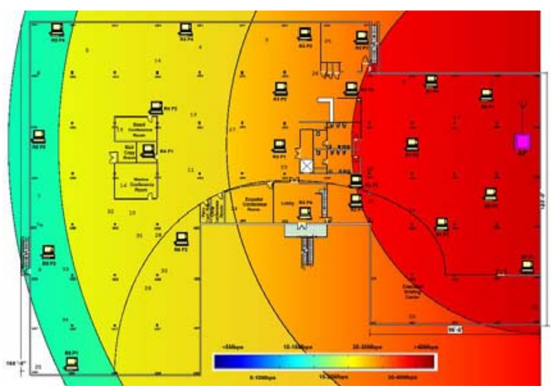

Figure 2 MIMO AP

The above testing shows the dramatic effect MIMO technology can have on a WLAN system, by providing much better throughput over the entire coverage area.

However, MIMO is not without its drawbacks, although there are relatively few when balanced against the gains. These are largely due to the increased component costs within an Access Point required for the MIMO implementation. Multiple antennas, and the use of higher performing CPU and DSP chips mean that, initially, manufacturer's costs of goods per AP may rise, however, as these components become mass produced it is likely that their prices will quickly fall. Additionally, the operation of multiple radio chains requires more power than the SISO (legacy) systems, both in transmit and receive operations. Clients that use standby mode to preserve battery life may have shorter cycle times, but work is already underway by client manufacturers to operate in single chain mode until a "wake-up" signal is detected.

The MIMO Architecture

Current WLAN systems use a Single Input, Single Output (SISO) architecture. While there is some spatial diversity in the availability of two receive antennas, only one antenna is actually connected to the receiver at any time; the receiver samples the signal strength at both antennas during the preamble and selects the best signal to use for collecting the data. Only one antenna is used to transmit. This technique is used to mitigate some multipath and signal fading in the environment.

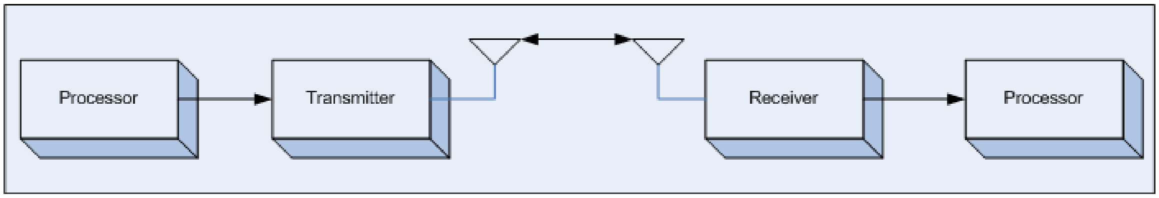

Figure 3 Existing SISO Architecture

In contrast, the MIMO system depends upon the receivers capturing a different combination of transmitted symbols and MIMO system throughput actually increases with the amount of multipath. Therefore, all receivers are used in a MIMO system to collect the multiple transmission streams and the DSP, added to the receive path, processes the received signals to separate the various transmitted streams and to recover the original bit stream. The DSP in the transmit path de-multiplexes the original bit stream according to the selected mode of operation. This can be to send multiple, independent bit streams (Space Division Multiplexing) to increase throughput or to code the same signal across the antennas (Space-Time Coding) to increase range and robustness of the signal.

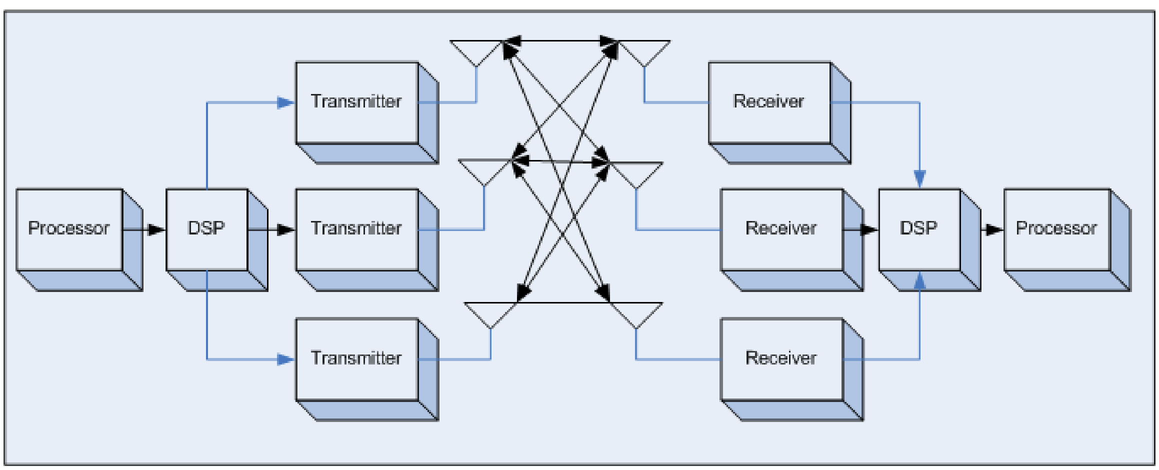

Figure 4 MIMO Architecture

System Implementations

Two general classes of system implementation are now possible. In a **Full MIMO** implementation both the AP and the clients employ MIMO technology to provide maximum throughput and coverage. A homogenous system can operate at maximum efficiency without the need to accommodate or protect for legacy systems.

The most common implementation is **Partial MIMO** which consists of a MIMO AP with a mix of MIMO and legacy (802.11a/b/g) clients. In this case the AP must provide for the legacy system operation of the clients. The result of this operational mode is system performance better than 802.11a/b/g. A similar effect is noticed when legacy 802.11b clients are mixed with 802.11g clients on an 802.11g AP.

The principle benefit to legacy clients in a Partial MIMO system is increased range. The ability of the AP to send the same bit stream a number of times in the same channel simultaneously affords the client the opportunity to pick the best available signal, and within the ability of the receiver, to use the combined signal strength to gain range, or to trade range for throughput. Increased range yields larger coverage areas per AP and reduces the deployment (CAPex) and operational (OPex) costs for the wireless network by reducing the number of APs to be deployed and the number of wired backhaul links to be installed.

A MIMO AP advertises its high throughput capability in the Beacon and Probe Response messages. A legacy client, in its Association Request, can only ask for basic 802.11a/g services since it has no MIMO capability. The AP, in granting access, marks the client as a legacy device and adjusts its (the AP's) mode of transmitting to and receiving from that particular client. Each client is treated independently so the network can handle a mix of MIMO and legacy clients. MIMO clients, on the other hand, seek and request the high throughput capabilities of the MIMO AP.

MIMO Examples

In this example, we will look at the benefits of MIMO over the current SISO-based APs as they manifest themselves in range and coverage. The table below contains the transmitter power, transmitter antenna gain and receiver sensitivity values for typical SISO and MIMO components. The processing advantages of MIMO will not be considered since only a single transmit/receive chain is being considered. For purposes of this example, a fade margin of 15 dBand a loss coefficient of 3, both suitable for indoor deployments, are used in the path loss and range calculations.

Parameter	SISO	MIMO
Tx Power	20 dB	14 dBm
Tx Antenna Gain	2 dBi	2 dBi
Rx Receiver Sensitivity	-70 dBm	-80 dBm
Fade Margin	15 dBm	15 dBm
Loss Coefficient	3	3
54 Mbps Range	17 m	23 m
54 Mbps Coverage	908 sqm	1662 sqm

Table 1 MIMO vs. SISO @ 54Mbps

The receiver sensitivity used is the value reported for 54 Mbps communication. The calculated 54 Mbps ranges for the SISO and MIMO systems are 17 meters and 23 meters, respectively. Thus the MIMO system has a 35% increase in operating range over the SISO system.

Using the calculated ranges, the coverage areas are 908 square meters (SISO) and 1662 square meters (MIMO). Thus the MIMO system has an 83% increased coverage area when compared to the SISO system.

If one examines the basic parameters, it is clear that the increased range and coverage are a result of the increased receiver sensitivity in the MIMO application. The transmit power of the MIMO equipment is 6 dBm lower than the SISO transmit power at the 54 Mbps data rate. This difference could be offset by using higher gain, say 6 dBi, omni antennas in the MIMO system, resulting in an even larger difference in range and coverage.

In the below example, we have reduced the data rate to 6 Mbps, the lowest IEEE 802.11g OFDM data rate; an 802.11g system can operate at lower data rates to be compatible with earlier 802.11b equipment. Because the 802.11g system uses a lower form of modulation at 6 Mbps than it does at 54 Mbps, it is able to both increase the transmit power, and to have better receive sensitivity.

Parameter	SISO	MIMO
Tx Power	20 dB	22 dBm
Tx Antenna Gain	2 dBi	2 dBi
Rx Receiver Sensitivity	-88 dBm	-95.5 dBm
Fade Margin	15 dBm	15 dBm
Loss Coefficient	3	3
54 Mbps Range	68 m	141 m
54 Mbps Coverage	14,527 sqm	62,458 sqm

Table 2 MIMO vs. SISO @ 6Mbps

The MIMO transmit power increases by 8 dBm and the receiver sensitivity moves to a value slightly above the noise floor in many locations. The SISO receive sensitivity also improves for the reason described above. The net result of these improvements is substantial increases in range for both systems, ~5-6X as shown above. The MIMO system retains a significant range advantage (107% or 2.1X) over the SISO system. The range difference, this time, results from improvements at both ends of the link.

As a result of the range improvements, both systems have significantly improved coverage areas. Again, the MIMO system retains a substantial advantage (329% or 4.3X) over the SISO system.

Test Results

Airgo Networks, a leading provider of high performance wireless technology, commissioned a study by the Tolly Group to evaluate the performance differences between MIMO and SISO systems. In the study, two APs and a client laptop using Airgo TrueMIMO™ radio modules were compared against standard SISO equipment for throughput and coverage in a representative office environment. Two objectives of the study were to determine the size of the Ethernet-speed (10 Mbps) coverage area and to create a throughput heat map showing throughput as a function of range from the AP.

The study has three principle results:

1. The MIMO systems demonstrated significantly better throughput over the SISO systems in both the 2.4 GHz ISM band and the 5 GHz UNII band. In the ISM band the MIMO throughput improvement was 10-fold over the SISO configuration while in the UNII band the improvement was 16 times. The improvement is manifested largely as a range-rate performance enhancement. MIMO systems typically have much better range-rate performance than SISO systems due to their ability to benefit from multipath, multiple independent signal streams and advanced signal processing.

2. As a result of the improved range-rate performance and the overall range improvements with MIMO, the MIMO coverage areas in both bands are substantially larger than the corresponding SISO coverage areas. Coverage area is determined as a circle, centered at the AP, with a radius equal to the maximum usable range. This range is generally associated with the lowest data rate at which the system operates. The MIMO coverage area in the ISM band was 19 times the size of the SISO coverage area; in the UNII band the MIMO improvement over SISO was 28 times.

3. The third principle result demonstrated the performance improvement of a mixed system in which a SISO client was associated with a MIMO AP. As explained previously, the MIMO AP is aware of the non-MIMO client capabilities and will transmit the same data stream from each transmitter. This enables the client to use the spatial diversity of its two antennas to pick the best signal of those sent. The improvements, while less dramatic than in a pure-MIMO implementation, are a five-fold improvement in throughput and a ten-fold improvement in coverage over a SISO-only implementation.

Multipath – Digging Deeper

Multipath is a common problem in wireless systems, particularly those intended to work indoors in cluttered environments. In the office or home environment, each vertical surface can reflect the radio signal causing a multipath signal to propagate within the environment. Multipath can manifest itself in both the analog and digital domains.

Analog multipath is most often seen as carrier interference at the receiver resulting in a null in the receiving pattern. Two or more carrier signals are received by the receiver and can enforce or cancel each other. The result is a rise (enforce) or drop (cancel) in the observed signal strength. In the worst case, the signals cancel each other completely and the receiver does not collect any data. This cancellation effect is most easily corrected by moving the receiver slightly to change the multipath situation. Since wavelengths for 2.4 GHz and 5 GHz WLAN systems are about 12 cm and 6 cm, respectively, moving the receiver several cm is all that is necessary.

Digital multipath results from the inability of the receiver to decode the data stream due to unrecoverable errors resulting from multiple, time-shifted, and otherwise distorted, copies of the data stream. In this case, more sophisticated receivers and additional digital processing are generally required to correct the potential problems. This effect is called inter-symbol interference because, as each symbol is being received, another version of the symbol arrives and modifies the version seen by the receiver.

Signal paths between a transmitter and receiver are typically some combination of direct and reflected. The direct signal path will always have the least loss because the path is shortest. Reflected signals will be delayed in time and attenuated due to the longer path and may suffer other distortions, e.g., inversion, due to the number of reflections and the reflecting surface composition.

The timing difference between signals reaching the receiver is called delay spread. Certain types of receivers, e.g., Rake, and signal modulations, e.g., OFDM, are more tolerant of delay spread than others. As the delay spread approaches a symbol time in length, the potential to confuse the receiver and create a false receive data stream is increased. Delay spreads in office environments may range beyond 100 nano-seconds because of the size and nature of the environment; home delay spreads are typically in the 10-20 nano-second range.

In SISO systems, multipath is a potentially harmful effect of the operating environment. In some cases, it is the only manner in which to have a client receive data from an AP as in the figure above. In this geometry, the AP and the client do not have a direct communications path due to the yellow obstacle. In this case, the receiver will attempt to recover the transmitted data stream for all signals arriving within its delay spread tolerance. If the Path 3 signal did not exist, the receiver would successfully recover the transmitted signal. However, the presence of the Path 3 signal within the time window of the receiver, results in unrecoverable errors in the receive process and the loss of the data.

MIMO systems, because of their inherent dependence on multipath and their use of multiple receivers, would collect the three sets of data and process them to recover the original signal. In fact, the three data sets would most like not be “copies” of the original data set, but subsets of the original data that can be mathematically processed to recover the original data stream. The ability to use signal characteristics and signal processing to make the communications more reliable is called processing gain.

Antennas

MIMO should not be confused with beam forming. In beam forming, the same signal is sent from multiple antennas, but is adjusted in phase and power to form a coherent signal in the direction of the desired client. Beam forming requires additional logic and firmware to properly implement and is probably unnecessary in a confined space such as indoor WLANs. The most likely uses for beam forming indoors would be a very high throughput application, or some specialized location service. Beam forming is more often used outdoors to form dynamic point-to-point links.

MIMO works extremely well with simple, omni antennas. The most often seen deployment is a small array of omni antennas, one per transmitter/receiver chain, either arranged linearly or in a triangular pattern. The effect of either arrangement is to essentially create an omni coverage pattern for the array. Since the MIMO system gains in the processing arena, the omni antenna array does not have any overall increased transmit or receive gain beyond the individual antenna gain specification.

Specialized antennas may be designed and deployed in unique circumstances. This may include coverage areas requiring directional antennas, or have specific deployment scenarios, e.g., all wall or all ceiling mount.

802.11n Implications

The 802.11 Working Group has created a Task Group “n” (TGn) to develop an amendment for the 802.11 Standard for WLAN equipment. This effort began in September 2003, following a study period to determine the feasibility of the task. The amendment will specify operation of the Medium Access Control (MAC) and Physical (PHY) layers of the network architecture; other IEEE 802 working groups are responsible for higher layers in the protocol stack.

802.11n requires the use of MIMO, but MIMO *does not* require 802.11n. This may seem counterintuitive, but the basic takeaway is that MIMO technology can ride over any radio technology; 802.11a, 802.11b, 802.11g and the future 802.11n, while 802.11n requires the use of MIMO technology to operate.

Some specific requirements were defined at the beginning of the task. These include;

1. At least one operational mode with data rates in excess of 100 Mbps in a 20 MHz channel as measured at the MAC Service Access Point (SAP). This a departure from earlier systems where the quoted data rate is the maximum bit rate in the air link. A data rate of 100 Mbps at the MAC SAP is roughly a four-fold improvement over existing 802.11 systems.
2. Support for the 5 GHz UNII (Unlicensed National Information Infrastructure) band, including Dynamic Frequency Selection, Transmit Power Control and operation in the 4.9 GHz Japanese band.
3. Some modes of operation that are backwards compatible with and interoperable with 802.11a and 802.11g.
4. Quality of Service (QoS) support through 802.11e.
5. To make efficient use of the available unlicensed spectrum, the highest throughput mode of the amendment must have a spectral efficiency of at least 3 bits/sec/Hz.

Unfortunately, 802.11n has been plagued by setbacks and amendments, and it is unlikely that we will see a ratified standard before late 2007.

Conclusions

MIMO technology is a significant improvement over the basic SISO systems that have been available up until now. As with any new technology, it will take some time before we see widespread deployments of client-side MIMO cards, but MIMO still provides over a 30 percent increased in throughput, range and data rates with legacy 802.11a/b/g clients as demonstrated in this paper. Perhaps most importantly, MIMO has been selected as the foundation for the new 802.11n standard, and while it will be at least another year before 802.11n is ratified, MIMO is clearly a long term technology that will be prevalent in all future access points.