ON Semiconductor

Is Now

To learn more about onsemi™, please visit our website at www.onsemi.com

Optimizing Wireless Resources In 802.11ac and Beyond

Originally published in Apr. 2014 under the company name Quantenna® Communications, Inc.

ON Semiconductor®

Optimizing Wireless Resources In 802.11ac and Beyond

Abstract

The extremely fast growing wireless data demand has fueled the interest in improving Wi−Fi® network capacity. Since spectrum resources are ultimately limited, serving more users and applications within existing resources has been the driving factor for research and innovation in this field. In this paper, we investigate how to exploit spectrum efficiency to address the needs of upcoming and future Wi−Fi networks. We present our view on the future Wi−Fi landscape with respect to market and technology trends. We then evaluate recent technology improvements and discuss upcoming advances to show the significance of these technologies for the future of Wi−Fi.

Introduction

New Wi−Fi Applications

Wi−Fi is becoming one of the most dominant interfaces ever. In 2012, 1.5 Billion new Wi−Fi enabled devices were sold, far exceeding the number of Ethernet enabled devices which was 1.2 Billion. It is projected there will be over 2 billion new Wi−Fi devices shipped in 2015. Furthermore, Wi−Fi is not only the dominant, but often the sole networking interface in many devices.

While originally intended for convenience and best effort connectivity, new applications and services together with Wi−Fi's market dominance are now forcing higher standards for Wi−Fi than originally envisioned. Further, the need and economics are forcing the de−commoditization of the WiFi.

The most dominant driver for internet traffic is video, which requires superb connectivity. It has been forecast that greater than 80% of the consumer traffic over internet will be video in 2016, further driven by the expected deployment of Ultra High Definition (UHD) or 4K TV today and the vision for 8k TV in the future. Wi−Fi is a typical mode of connectivity for OTT video services.

Consumer video−centric devices such as TVs, gaming stations, and media centers are poised to demand better Wi−Fi connectivity as well.

Popularity of other OTT applications such as voice over Wi−Fi is also increasing rapidly. Voice over Wi−Fi has become so popular to the extent that a number of new Mobile Network Operators (MNO) are using Wi−Fi as the primary connection and cellular as a fallback.

Another fast growing internet application is “cloud”, defined as “dumb” display devices enabled by connecting to storage and intelligence residing elsewhere. Connectivity is paramount for these devices.

Internet of Things (IoT) is another new and fast growing field requiring wireless connectivity. As IoT matures, the number of wireless devices will surge. With that, Wi−Fi will be a critical for applications like smart home.

Outdoor Wi−Fi is becoming a way to provide broadband internet service. Wi−Fi can be used to provide hundreds of megabits per second over thousands of feet under certain outdoor deployment cases. This is considered a very cost optimized way to provide access for locations where wired access is not available or is too challenging and costly. Wi−Fi is also very critical in enterprise environment and education industry. Schools are adopting Wi−Fi at a fast pace and companies are replacing desktops with laptops. Reliable Wi−Fi is becoming essential in all of these environments.

For internet service providers, serving multiple video and data clients inside the home requires high Wi−Fi throughput, coverage and reliability. Wi−Fi is becoming an attractive option for service providers because it allows them to better control their customer’s internet service experience. In addition, Wi−Fi reduces deployment costs by eliminating in−home wiring requirements for paid video distribution.

Carrier Wi−Fi is another growing field among internet service provider. Consumers like to have Wi−Fi access in more places outside of their home or office. Service providers satisfy this need by offering Wi−Fi access through wireless hotspots. As a result, the number of carriers owned Wi−Fi hotspots is growing very rapidly.

Wi−Fi is becoming an important part of strategic planning for mobile network operators as well. Handling the exponential traffic growth in mobile networks requires Wi−Fi traffic offloading. This has resulted in significant research and development in the area of Wi−Fi small cells and heterogeneous networks among MNOs.

All previously mentioned applications and use cases depend on reliability, range and speed of Wi−Fi networks.

Asymmetric APs and Clients

Access Points (AP) are the heart of Wi−Fi network. Network performance is highly dependent on AP capabilities. Therefore, having an advanced AP in the network which can reliably serve many clients is very important.

There is a shift in industry toward high−end APs. These APs support a large set of optional standard capabilities such as high order of MIMO, beamforming and Multi−User MIMO (MU−MIMO). These devices are capable of delivering data rates in excess of a gigabit per second.

However, on the client side, space, cost and processor limitations do not allow for implementation of such advanced configurations. As a result, divergence between AP and client capabilities is happening.

This brings a lot of new challenges and opportunities into Wi−Fi landscape specifically in regards to spectral efficiency.

There are a number of different ways to boost wireless capacity: increasing the bandwidth, power or improving the spectral efficiency. During evolution of Wi−Fi in the last two decades, all these options have been aggressively exercised.

While increasing the bandwidth has proven to be relatively straightforward, the capacity improvement gained by improving spectral efficiency is much more valuable because spectrum is extremely precious. With limited bandwidth and power, increasing efficiency is critical.

Multiple Input Multiple Output (MIMO), higher modulation and advanced coding schemes can all help in this, at the expense of more sophisticated (and hence, costly) Silicon. However, silicon costs are trending downward making the investment in these techniques the better option.

A survey of silicon production cost over next decade shows that silicon prices are expected to decrease significantly. Figure 1 shows these results for 28, 40 and 65 nm.

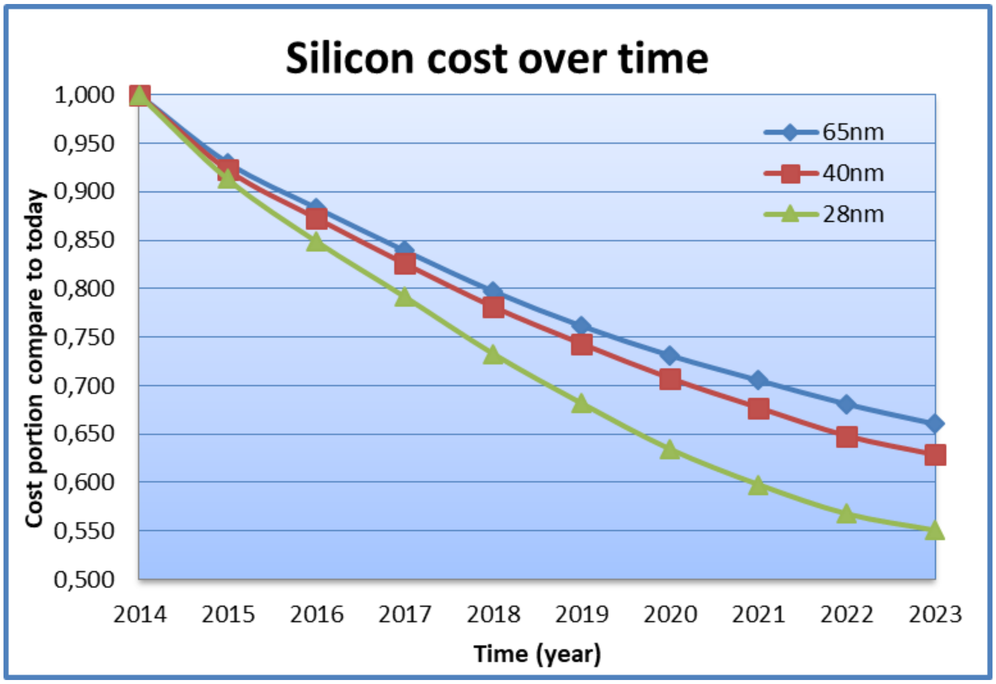

Figure 1. Estimated Silicon Cost over Time over Next Decade

Therefore, it is predicted that Wi−Fi industry will shift more toward investing in advanced solutions which will be more spectral efficient. Many of the changes have been already adopted by the IEEE 802.11 standard.

In this paper, we first review the essentials of the MIMO channel. We then highlight how using MIMO contributes to better robustness and efficiency. After that, we elaborate on the expected direction of the Wi−Fi technology. We consider advanced MIMO techniques and other approaches that will play a role in future of Wi−Fi.

Understanding the MIMO Channel

The term “MIMO” refers to communication systems with multiple transmit and receive antennas. The transmit signal consists of N_{TX} components (one for each transmit antenna), while the received signal consists of N_{RX} components (one for each receive antenna). Mathematically, transmit and receive signal can be represented as N_{TX}−dimensional and N_{RX} −dimensional vectors respectively. The channel between transmit and receive antennas can be represented as a linear transform that converts the transmit vector into the receive vector. Such an operation mathematically corresponds to matrix multiplication where the dimensions of the matrix are N_{RX} x N_{TX}. This equivalence is illustrated for a 2x2 MIMO system in Figure 2.

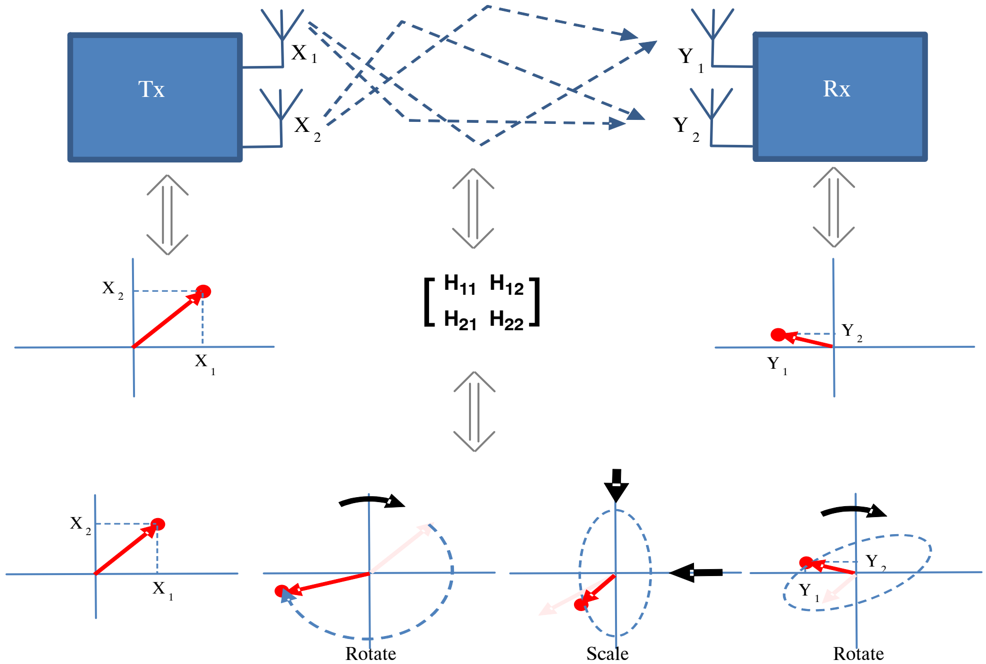

Figure 2. MIMO Elements and Their Mathematical Representations. MIMO Signals Become Vectors. The MIMO Channel Becomes a Matrix

In two dimensions, the transmitted and received signals both consist of two independent components. We can represent these two−dimensional signals as points in a plane, with one of its components corresponding to the x−axis and the other to the y−axis. The 2x2 channel is represented by a 2x2 matrix that transforms the transmit point (X_1,X_2) into the receive point (Y_1,Y_2).

A fundamental result of matrix theory states that any matrix can be decomposed using Singular Value Decomposition (SVD) into a sequence of three basic operations. Specifically, the operation of the matrix on a vector is equivalent to the following series of operations:

1. Rotate the vector
2. Scale the Components of the Rotated Vector
3. Rotate the Scaled Vector

This is illustrated in as part of Figure 2.

Understanding different intermediate steps is the key to understanding some of the properties of MIMO channels. Specifically, the scaling operation may cause a degradation of the signal quality, since some signals may be “badly oriented” and undergo significant attenuation as a result of scaling.

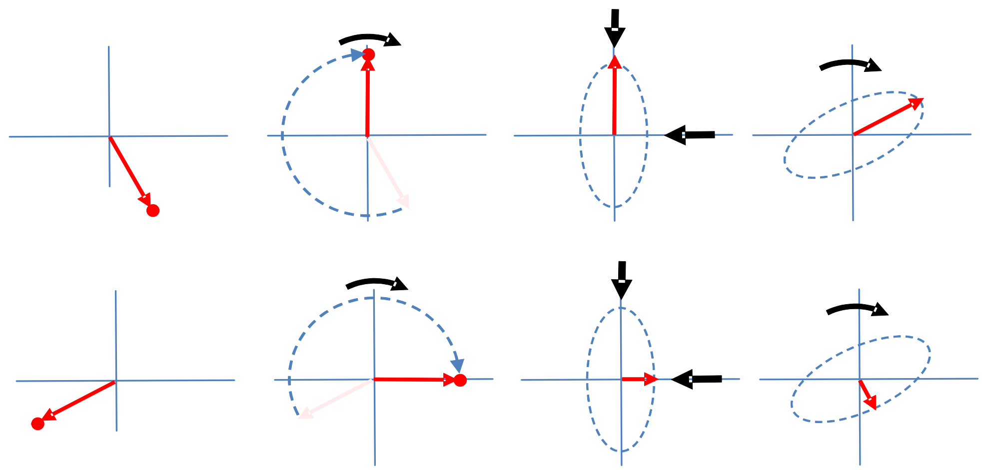

Figure 3. Two Transmit Signals with Equal Power Traveling Through the Same Channel Result in Very Different Receive Signals

Figure 3 shows two transmit signals (the upper and lower part of the figure), each represented as a two−dimensional point. Both transmit signals are sent through the *same* channel, meaning both undergo the same series of rotations and scaling.

For the first transmit signal, initial rotation results in an intermediate point that lies along the vertical axis. In contrast, the first rotation of the second transmit signal lies along the horizontal axis. For this particular channel, the horizontal axis undergoes a much larger attenuation than the vertical axis. As a result the second signal is much weaker at the receiver than the first signal. It is important to realize that the only difference between the two transmit signals is a rotation. Both signals have the same power, but they are distributed differently over the two transmit antennas.

SVD explains why signals with the same transmit power can lead to very different receive signals even as they travel through the exact same MIMO channel. Understanding the reasons behind this will allow us to optimize signal reception. From Figure 3, one can see that a simple rotation of the “bad” signal would have turned it into the “good” signal achieving optimal receive power. Channel information is required to determine the rotation (or “precoding”) that should be applied.

Technically, the first rotation is described by the so−called “V−matrix”, which is central in the channel feedback provided by an 802.11ac receiver. The V−matrix alone (i.e. the first rotation) can be used to determine the pre−rotation that should be applied to achieve the maximal received signal.

The scaling operation can also reduce some of the components of the signal to zero. This means that it is possible to transmit a signal through a MIMO channel and receive zero signals at a particular receiver. This property becomes very useful when considering Multi−User MIMO (MU−MIMO) which we will discuss later.

MIMO Concepts

In this section, we'll discuss some important MIMO concepts and highlight how MIMO can help in providing more robustness or higher throughput.

Transmit Diversity

Transmit diversity refers to a case where correlated signals are transmitted from multiple antennas. In its simplest form, the same signal could be replicated across the antennas. Obviously, the total transmit signal power has to remain constant, so the more antennas, the less power can be transmitted on any individual antenna. For a system with N_{TX} transmit antennas, each antenna will roughly carry 1/ N_{TX} of the total power.

The receiver will now observe the sum of N_{TX} replicas of the transmitted signal. Each replica follows a slightly different path from transmitter to receiver. As a result, the received signal is not victim to fading in a single path. The signal replicas average out at the receiver and tend to smooth out any irregularities in the medium.

Figure 4(a) shows that adding transmit antennas with transmit diversity alone does not necessarily increase performance, but it does improve the robustness in a fading channel environment.

Receiver Diversity

Receiver diversity refers to the case where the receiver employs multiple antennas to receive the signal and then combines the signals on the different antennas for optimal performance. Unlike the case with transmitter diversity, each additional antenna does increase the total received power. For instance, increasing antennas from one to two, leads to doubling of the received power. It is up to the receiver to combine the signals received on the various antennas to optimize the performance. In Maximal−Ratio Combining (MRC), the different signals are added together with a weighting that is proportional to the signal strength of that path. Since throughput increases with higher received power, the performance increases with any increase in the number of receive antennas.

Figure 4(b) show performance increases with the number of receive antennas.

Spatial Multiplexing

The real power of MIMO systems comes from spatial multiplexing. When both transmitter and receiver have multiple antennas, the link between transmitter and receiver can be thought of as consisting of multiple independent transmission paths. For a system with N_{TX} transmit and N_{RX} receive antennas, the number of independent paths is equal to the minimum of N_{TX} and N_{RX}.

In such an environment, a MIMO system can transmit multiple independent data streams over the existing transmission paths. This is known as Spatial Diversity Multiplexing (SDM) or Spatial Multiplexing (SM).

Importantly, this improvement in data throughput comes without any increase in transmit power or signal bandwidth. Effectively, MIMO combined with spatial multiplexing opens up another dimension that can be used for increased performance.

Figure 4(c) shows how spatial multiplexing increases performance. As the MIMO order grows, the throughput increases steadily.

Transmit Beamforming

While MIMO can create multiple independent paths between the transmitter and the receiver, not all paths are of the same quality. If the transmitter has access to channel information, it can precode its signal to choose the best transmission modes. This is what transmit beamforming does.

Precoding means that signals are sent with specifically chosen amplitudes and phases over the available antennas. This amounts to a weighting of the signals with complex coefficients. Since the optimal weighting depends on the channel between transmitter and receiver, some form of channel information is required. This channel information can be supplied through a direct feedback mechanism from the receiver or derived by the transmitter using alternative means.

Figure 4(d) shows the performance gain with and without beamforming for a number of different MIMO configurations. This figure shows that Beamforming can provide significant performance improvement for transmitters with high number of antennas.

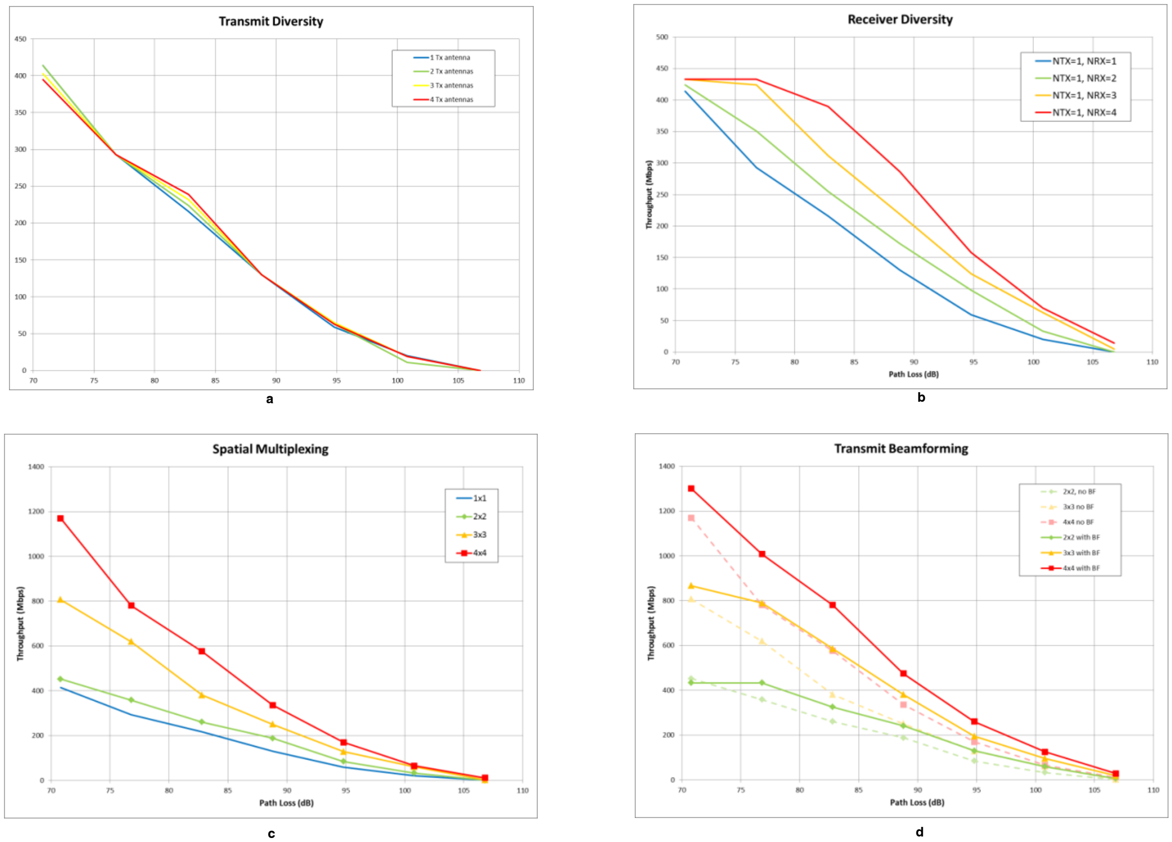

Figure 4. Performance Impact of Various MIMO Configurations

802.11ac

802.11ac is the name of one of the latest additions to the successful family of 802.11 wireless protocols. 802.11ac was officially published in December 2013 as Amendment 4 to the IEEE 802.11−2012 standard.

802.11ac operates strictly in the 5GHz band, but supports backwards compatibility with other 802.11 technologies operating in the same band (most notably 802.11n).

To achieve its goal of true wireless Gigabit, 802.11ac relies on a number of improvements in both the MAC and Physical Layer (PHY). 802.11ac also shares many features with 802.11n. Advanced coding (LDPC) is used for increased coding gain. STBC technology can be used for transmitter diversity. A number of other features provide various levels of performance improvement, such as explicit feedback Transmit Beamforming, and short GI.

The most notable feature of 802.11ac is the extended bandwidth of the wireless channels. 802.11ac mandates support of 20, 40 and 80 MHz channels. Optionally, the use of contiguous 160 MHz channels or non−contiguous 80+80 MHz channels is also allowed.

802.11ac optionally allows the use of 256 QAM in addition to the mandatory BPSK, QPSK, 16 QAM and 64 QAM modulations. 256 QAM increases the number of bits per sub−carrier from 6 to 8, resulting in a 33% increase in PHY rate under the right conditions.

The 802.11ac standard uses the concept of “spatial streams” to designate independent data streams that can be transmitted using spatial multiplexing. The maximum number of spatial streams is increased from four to eight. Devices with 8x8 MIMO configration will harvest the maximum performance in 802.11ac standards.

802.11ac introduces a single transmit beamforming feedback protocol, streamlining the various different beamforming modes that were originally included in 802.11n. Channel information is obtained by sending special “sounding frames” or “Null Data Packets” (NDP). This allows the receiver to estimate the channel. The receiver then performs a Singular Value Decomposition (SVD) on the estimated channel matrices and sends back partial information from this decomposition (the so−called V−matrix) for all or for a subset of the tones.

Downlink MU−MIMO was added to 802.11ac to address the multi−STA throughput requirement. In MU−MIMO, the Access Point (AP) transmits independent data streams to several STAs at the same time. Through preprocessing of the data streams at the transmitter (similar to what happens in beamforming), the interference from streams that are not intended for a particular STA can be eliminated. 802.11ac accommodates a maximum of four simultaneous users per transmission. Each user in a multi−user group can receive up to four streams, with the total number of streams not exceeding eight.

Frame aggregation was introduced first in 802.11n to increase the transmission efficiency. 802.11ac takes this concept further by increasing the maximum frame aggregation size from 65,535 octets to 1,048,575 octets.

The 802.11 RTS/CTS mechanism has been updated to support a “dynamic bandwidth” mode. The new mechanism allows devices to determine whether all or part of the channel is available. 802.11ac can use this information to fall back to lower bandwidth transmission. This helps mitigate the effect of hidden nodes and optimizes spectrum use.

The Future Wi−Fi Landscape

APs vs. Client Devices

In the Wi−Fi consumer space, a clear divergence can be observed between the capabilities of APs and STAs. Typically, the AP is a stationary device that is connected to the power grid. STAs like smartphones and tablets on the other hand tend to be more and more mobile and battery−powered. This trend is likely to continue in the future. Because of practical constraints, these Wi−Fi STAs will be limited in the number of antennas they can support and the amount of processing the can deliver.

We have to recognize and incorporate this asymmetry between APs and STAs to be able to deliver the best performance in this environment. Therefore, the focus should be to equip both AP and STA with capabilities that enable the end−to−end system to get optimal throughput.

Below, we will highlight some of the techniques that are beneficial in this environment.

Beamforming and MU−MIMO

One of the appealing aspects of beamforming is that it allows for performance gains while only imposing a minimal complexity penalty on the receiving device. This is because the onus of applying the precoding is mostly on the AP (including the fact that multiple antennas are needed). The complexity of the AP obviously increases when transmit antennas are added. However, the complexity of the client device is only minimally affected. The only requirement on the client is that it provides channel feedback that can be used by the AP.

Likewise, MU−MIMO requirements for client devices are fairly minimal, but an MU−MIMO enabled AP can potentially increase the total delivered downlink throughput by multiple folds. In MU−MIMO, the AP sends two or more simultaneous signals which are superimposed on each other and each signal is intended for a specific receiver. Ordinarily, simultaneous presence of multiple signals would appear as interference, but in MU−MIMO signals are precoded at the transmitter such that the signal is received free of interference at each receiver. The key to achieve this is to exploit zero modes in channels from transmitter to the various receivers. This can again be illustrated in two dimensions to give a more intuitive understanding of the process.

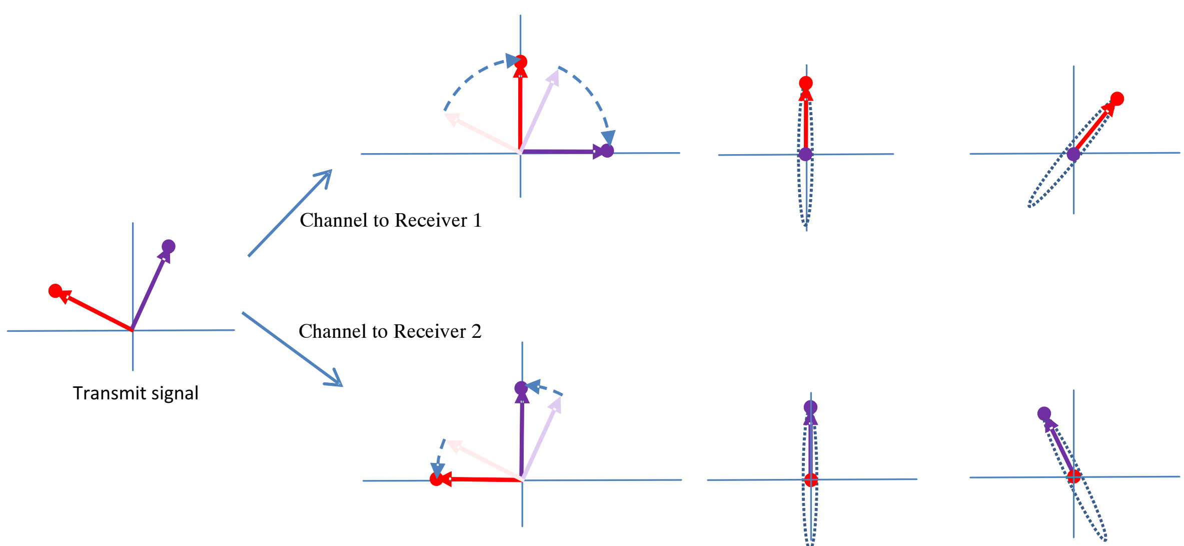

Figure 5. MU−MIMO Ilustration

In Figure 5, we consider a case where the AP addresses two users with a MU−MIMO transmission. The transmit signal consists of two sub−signals which are sent out simultaneously. The combined signal travels towards the first and second receiver over different channels. In Figure 5, channels have the special property that the signal component along the horizontal axis becomes zero. MU−MIMO achieves interference−free communication when the AP manages to find a precoding that maps the sub−signal for the second receiver into the null space of the first receiver and the sub−signal for the first receiver into the null space of the second receiver. As can be seen from the simple illustration in Figure 5, each receiver gets an interference−free signal that only contains the sub−signal that was intended for that receiver.

To guarantee the existence of zero modes, the total number of receive antennas should not exceed the number of transmit antennas. Therefore, increasing the number of antennas at transmitter provides more degrees of freedom that can be optimized to discover mutual zero modes between devices while still optimizing performance of each link. Also, more antennas mean more devices can be addressed simultaneously, making sure that no channel capacity goes to waste.

Figure 6 shows some results for MU−MIMO. Here, each client device is a single−stream device. This would normally limit throughput at low path loss. By using MU−MIMO, channel capacity is exploited much more efficiently.

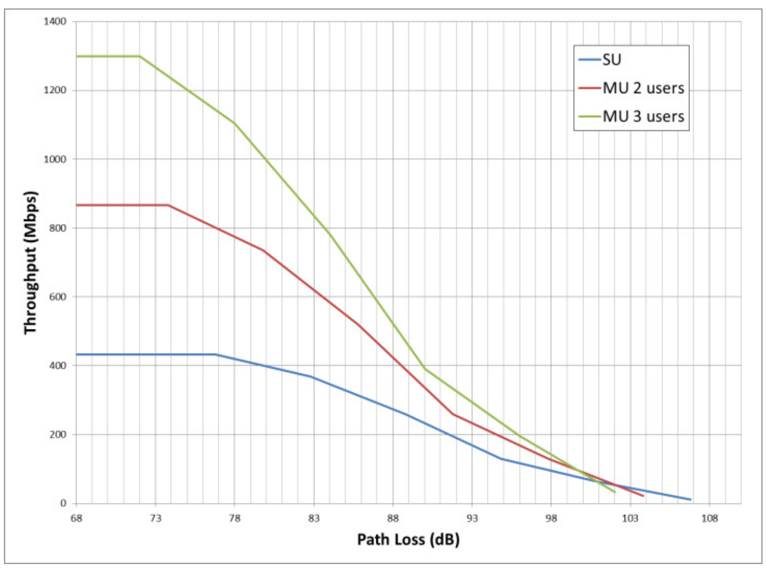

Figure 6. MU−MIMO Increases Total Downlink Throughput

UL−MU−MIMO

Even though not part of the 802.11ac standard, Uplink MU−MIMO (UL−MU−MIMO) can be an enhancement for future WiFi devices. In uplink MU−MIMO, multiple transmitters simultaneously send data to a single receiver. In the context of Wi−Fi, this would correspond to a case of multiple client devices communicating with an AP. The receiver can distinguish between the different users if it has a sufficient number of antennas and if the users are sufficiently spatially uncorrelated. Using techniques like successive interference cancellation (SIC), the AP can recover the data from the various users.

Other forms of MU−MIMO that are not currently part of 802.11ac may also gain importance in the future. For instance, 802.11ac only allows the AP to transmit MU−MIMO packets. This could be extended to allow any station to send point−to−multipoint traffic using Multi−user techniques.

Orthogonal Frequency Division Multiple Access (OFDMA)

OFDMA is a combination of time and frequency division access. OFDMA is a concurrent multi−user access mechanism achieved by allocating partitions of time−frequency resources to each client. OFDMA has not been considered in WiFi yet but is used in other wireless technologies such as LTE and WiMAX.

The major problem with frequency division approaches is interference among frequencies which requires higher performing filters. On the other hand, FDMA does not have access timing issues which come with TDMA. It can also provide lower latency since it can provide simultaneous continuous channel access to many users.

Considering higher bandwidth possibilities of 802.11ac, potential for supporting frequency division access in WiFi has increased. While high bandwidth access points are entering the market, still many clients will only support a subset of this bandwidth. This means that for users with these legacy devices, network efficiency will stay low even if they upgrade to higher bandwidth access point.

With OFDMA, this inefficiency can be improved. A 160 MHz bandwidth access point can simultaneously talk to four 40 MHz devices in different frequency. This means network efficiency can be increased up to 400% compare to current access mechanism. In addition to higher network efficiency and better aggregate throughput, other QoS parameters like latency and packet error rate can also be improved with OFDMA. The frequency separation reduces the collision possibility and allows lower latency by providing concurrent channel access for many users. With OFDMA resource allocation granularity can be also increased due to existence of two dimensions for resource allocation.

Full Duplex Operation

One interesting, but challenging, new technology is full duplex operation. Currently, transmission and reception alternatingly use the same frequency band. Full duplex operation refers to a mode of operation where the same station simultaneously transmits a signal and receives another signal. Because of the full spectral overlap, significant effort is needed to separate the two signals such that the (strong) transmit signal on the antennas does not drown out the (much weaker) receive signal on the same antennas.

A lot of progress has been made in this area for single−antenna systems. The most advanced technologies, which rely on a combination of analog and digital separation, can achieve signal rejection of up to 110 dB. Unfortunately, for MIMO systems a lot of challenges remain to be addressed.

In addition to changes that are needed in the PHY, the Wi−Fi MAC would also have to be modified to accommodate this new mode of operation.

Ideally, the AP would be able to interact with multiple low−end clients (e.g. transmitting to one while receiving from the other) without these clients having to incur extra complexity. The burden for making simultaneous transmission and reception possible would again fall on the AP.

Given the remaining challenges for full duplex operation, this is still an active area of research. However, the potential doubling of spectral efficiency that comes with full duplex operation makes it an exciting option for future Wi−Fi systems.

Self−Organizing Networks

Another growing field in Wi−Fi landscape is network management. Wi−Fi networks traditionally have been managed statically. In static management, devices were programmed with fixed configuration parameters, which could result in huge network inefficiency. For example, an AP located in a large house should use high transmit power for better coverage. Meanwhile, when similar AP is located in a small apartment, it can provide full coverage and avoid interference on neighboring apartment networks by using a lower transmit power. Therefore, having same transmit power for these two APs is not an optimized solution and can result in lower network efficiency. Similar improvements can be done for other parameters such as operating frequency or receive sensitivity.

As a result, Self−Organizing Network (SON) network concept emerged where vendors added dynamical algorithms to adjust these parameters depending on deployment environment. In many new products, each device by itself evaluates its statistics and decides on its configuration parameters. This is called Distributed−SON (D−SON). But since each device has access to limited information, these decisions may not be globally optimized. Therefore, other centrally managed solutions called Centralized−SON (C−SON) are arising to improve efficiency. C−SON has its own problems such as reaction time and high overhead. Network performance can be impacted due to required communication between each device and management entity. As a result, Hybrid SON solutions are offered to provide a balance between D−SON and C−SON.

This is still a growing field and it is expected that new solutions in this area to emerge in future to better address network efficiency issues related to cross network optimization.

Conclusion

To satisfy the ever−increasing demand for Wi−Fi throughput, new wireless devices will have to rely on emerging and to−be−developed techniques to use the available wireless resources to the maximum. Central piece of this approach is spectral efficiency that MIMO can provide. Many of these techniques like Beamforming and MU−MIMO are available today or will be in next−generation products. In addition, several future improvements may help to further address the needs of Wi−Fi deployments.

802.11ac, the latest addition to the IEEE 802.11 standard, already equips us with some of the needed methods. In addition, IEEE is starting a next−generation project that encompasses a list of possible improvements that are still in their initial stages today. Together, techniques discussed in this paper will ensure that Wi−Fi can live up to its promises to satisfy demand for new generation of emerging applications.

References

[1] IEEE 802.11−2012, "IEEE Standard for Information technology − Telecommunications and information exchange between systems, Local and metropolitan area networks, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", March 2012.

[2] IEEE 802.11ac−2013, "IEEE Standard for Information technology − Telecommunications and information exchange between systems, Local and metropolitan area networks, Part 11, Amendment 4: Enhancements for Very High Throughput for Operation in Bands below 6 GHz", December 2013.

[3] O. Bejarano, E.W. Knightly and M. Park, "IEEE 802.11ac: From Channelization to Multi−User MIMO", IEEE Communications Magazine, Vol. 51, Issue 10, Oct 2013, pp. 84–90.

PUBLICATION ORDERING INFORMATION

LITERATURE FULFILLMENT:
Email Requests to: orderlit@onsemi.com

ON Semiconductor Website: www.onsemi.com

TECHNICAL SUPPORT
North American Technical Support:
Voice Mail: 1 800−282−9855 Toll Free USA/Canada
Phone: 011 421 33 790 2910

Europe, Middle East and Africa Technical Support:
Phone: 00421 33 790 2910
For additional information, please contact your local Sales Representative