THE THINGS THEY DON'T TELL YOU ABOUT 802.11AC

THE PROMISES AND LIMITATIONS OF THE NEW IEEE WI-FI STANDARD

The Super Gigabit Wi-Fi standard - 802.11ac – is being widely adopted across the industry. By implementing new technologies and expanding existing capabilities, this feature packed standard promises to deliver a huge performance boost. But, in reality, there are some serious limitations that drastically affect the 802.11ac standard's ability to deliver on its promises. Read on to find out all about them.

THE THINGS THEY DON'T TELL YOU ABOUT 802.11AC

THE PROMISES AND LIMITATIONS OF THE NEW IEEE WI-FI STANDARD

"IN ALL AFFAIRS, IT'S A HEALTHY THING NOW AND THEN TO HANG A QUESTION MARK ON THE THINGS YOU HAVE LONG TAKEN FOR GRANTED."

BERTRAND RUSSELL

THE 802.11AC PROMISE

The new 802.11ac standard is being heralded as the 'Super Gigabit Wi-Fi' across the industry. It promises to take something great and make it even better. 802.11ac is a faster and more scalable version of 802.11n; 802.11ac couples the freedom of wireless with the capabilities of Gigabit Ethernet.

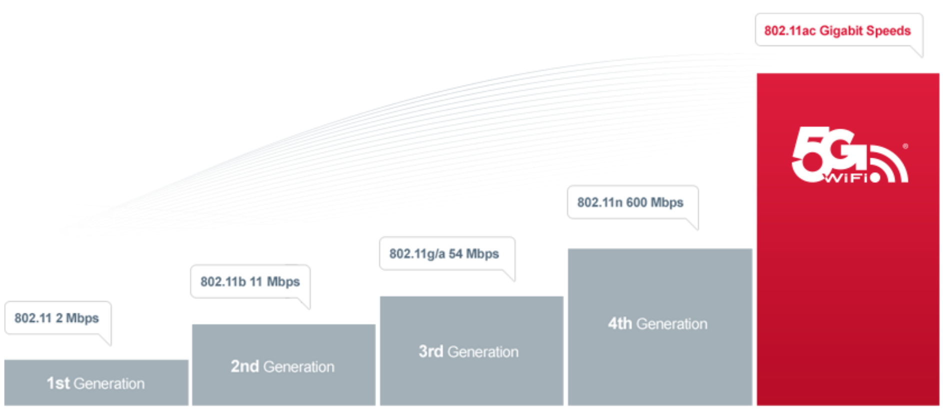

Everywhere you read the 802.11ac standard sounds like the movie trailer of a sure box-office hit. The crowds are gathering; Apple, Samsung and HTC smartphones already support first wave chips, as do laptops and tablets. This trend is expected to continue making

802.11ac available on seven of 10 smartphones in 2015 according to ABI Research.

In the 2nd half of 2014, the 2nd wave of 802.11ac chips will hit the market. These chips will support the much awaited multi-user feature known as MU-MIMO. With this feature, 802.11ac is expected to deliver even more, to more users.

The new 802.11ac standard incorporates some very good enhancements to the current technology. However, one needs to be careful and see that beyond the initial rosy picture, there are some critical factors limiting the 802.11ac performance in the 'real-world'. In the following chapters we will review the promise and limitations of each of these items in greater detail.

THE CAPACITY BOOST

Let's start by taking a closer look what makes the 802.11ac the 'Super Wi-Fi'. 802.11ac achieves its raw speed increase by pushing on three different dimensions:

1. Increasing the bandwidth from the maximum of 40 MHz in 802.11n up to 80 MHz or even 160 MHz in 802.11ac (for 117% or 333% speed-ups, respectively).
2. Higher modulation, now using optional 256 quadrature amplitude modulation (QAM), up from 802.11n's maximal 64QAM (for a 33% speed burst at shorter, yet still usable, ranges)
3. More multiple input, multiple output (MIMO) spatial streams (SS). Whereas 802.11n stopped at four spatial streams, 802.11ac goes all the way to eight (for another 100% speed-up). (Note that in practice 11n chipsets supported 3SS while 11ac wave 2 chips will support 4SS)

Together, these 3 dimensions translate to a maximal ***338%*** rate increase in the coming wave II chipsets (and an amazing theoretical maximum of ***1380%*** with the optional 160 MHz).

What they forget to tell you about the capacity boost

MAC INEFFICIENCY

The degradation in MAC efficiency is a small fact that is usually left unspoken. In the 'old' 802.11n standard operating at 20 MHz the MAC operated with over 85% efficiency. By comparison, the 802.11ac MAC efficiency is 65%. We are 'wasting' 20% more time on overheads.

> ***Relativity Theory***
>
> MAC Efficiency = Net data air time Relative to the air time used including all the overhead.

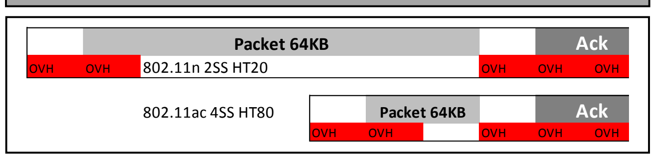

The reason for this change is the 'Relativity Theory' (it's a simple one, don't worry). In 802.11ac, the data rates increased significantly. As a result, the data now takes less air time. However, everything else is transmitted at the same old slow pace to ensure interoperability with older device that operate on 802.11a/n. This means that the - preambles, ACK packets, RTS, CTS and the CSMA protocol - take up just as long as before. Since the data now takes less time while the rest remains the same, the data now takes up only 65% of the stream in comparison to the older HT20 802.11n, where the efficiency was above 85%.

To make matters worse, there is no provision in the 802.11ac standard (unlike the 802.11n) for a Greenfield deployment. So, even in a theoretical deployment, where all the clients support 802.11ac, the non-data speeds will remain the same.

The implications of this lower efficiency reduce the enhancement dramatically from its promised 338% to less than 250%. While this still sounds great, take into account that we now use twice the spectrum. More importantly, in most cases much of this improvement will not be available at all (256QAM and 4SS utilization will be rare and require a 4 antenna client in proximity to AP). This leaves us with results that are similar to using 2 11n APs with 2 HT40 channels. In the next chapter we will further explain why this double bandwidth advantage may not deliver twice the capacity.

INSENSITIVITY TO NOISE

Wider bandwidth increases the capacity – in ideal conditions. In the real world, the noisy conditions pose serious limitation on the performance, especially in the case of Wi-Fi where different channels experience different noise levels.

Here is why.

In order to adopt wider bandwidths, wide band DMT/OFDM technologies choose a specific constellation per tone/subcarrier instead of using same constellation for all tones, since the SNR varies significantly between tones. In Wi-Fi, however, the constellation does not change over the entire bandwidth. As we move up to 80 MHz, the probability of noise diversity between 20Mhz channels and tones increases significantly. To limit errors and keep the Packet Error Rate (PER) above its required 90%, the rate control will cut back reducing the constellation and transmission rate. Since the same constellation / number of bits is encoded by all tones, the worst performing tone will determine the constellation for all the tones.

Let's use an example to highlight the effect of noise on the effective bitrate.

Scene 1: We have 2 40 MHz channels. The first 40MHz is operating at full 11n speed of 450Mbps. The 2nd one suffers from noise coming from a neighboring network and is forced to cut back to 117 Mbps.

Scene 2: enter 802.11ac into the arena. Now we have 1 channel that operates at 80MHz to replace our 2 40 MHz channels. However, due to our noisy neighbors, the channel has to throttle back to 234Mbps.

In theory, the total bitrate of the two scenes should have been equal. In practice, the high rates may actually achieve the opposite in many 'real world' situations.

LOWER SIGNAL TO NOISE

By moving to 80 MHz we lose twice reducing the Signal to Noise Ratio (SNR) significantly. On the S side of the SNR, when the AP transmits on 80 MHz, the power is stretched 2-4 times over the wider bandwidth. As a result, the power per hertz or per tone drops by 3 to 6dB. In addition, on the client side, the receiver is open to the full 80Mhz making, which increase the total noise by at least 3 to 6dB. The total SNR is degraded on both sides reducing the possible constellation, which reduces the Wi-Fi rates and coverage radius.

THE TRANSMIT BEAMFORMING

Beamforming was first introduced into the 802.11 standards as an optional and non-standardized element of the 802.11n spec. In the 802.11ac standard, beamforming will remain an optional feature, but its implementation will be standardized and more widely adopted as a result.

Omnidirectional vs. Beamforming

Traditionally, access points have been equipped with omnidirectional antennas, which are so called due to their 'ball' like transmission pattern. An alternative method of transmission is to focus energy toward a station, a process called beamforming. In beamforming the energy is focused toward a client, and is constantly moving per packet between the clients. By sending radio energy in one direction, it is possible to reach farther and achieve much better SNR to deliver higher capacity.

The IEEE 802.11ac standard supports only explicit beamforming. In the previous 802.11n standard, implicit beamforming was supported as well.

Implicit beamforming relies on the Access Point to analyze the channel by itself (by analyzing received packets). To achieve this; the AP will need to calibrate itself with a high degree of accuracy of the power levels, phases for both downlink and uplink from each antenna. This requires frequent and very complex internal measurements since the phases and power levels change constantly over time and temperature variations. Due to the complexity, the 802.11n implicit beamforming option was never used in standard off the shelf products.

Explicit beamforming is much simpler to implement and it is the option supported by the 802.11ac standard. In explicit beamforming there is a calibration session between the AP and client prior to using beamforming to transmit data.

To implement explicit beamforming, the 802.11ac standard defines a training sequence:

1. The training signal is transmitted from the AP to the client
2. The client performs channel estimation on this training signal
3. The client sends its results within a standard 802.11 packet back to the AP.
4. This feedback is called sounding in the IEEE standard and commonly referred to as CSIT.

The 802.11ac beamforming improves the performance of the underlying protocol by 2 to 6 dB. However, before rushing to deploy the 802.11ac beamforming, there are several things one may want to know about it.

What they usually forget to tell you about explicit TxBF

EXTEND THE RANGE WITHIN RANGE

One of the benefits of beamforming is extended range. However, during the training phase, the AP operates in a non-beamforming mode. It only switches to beamforming after the calibration has been completed. To participate in the training session, the client has to be located within the shorter 'non-beamforming' range. Furthermore, due to the dynamic conditions, the training session has to be repeated periodically. This means that explicit beamforming cannot extend the range.

COMMUNICATION OVERHEAD

The repeated training and sounding information imposes an additional communication overhead. This means that the AP is limited by the amount of users that it can support before most of its air time is spent on training and not actual data. Even with only a few clients, it is a delicate balance between the beamforming gain and bandwidth wasted on the training overhead. In tranquil residential areas, the beamforming may be beneficial. However it will be unusable in larger-scale carrier and enterprise environments, where the dynamic environment requires sounding every few milliseconds to update the channel estimations. The IEEE channel models for indoor deployments suggests that each client will need to repeat the training sequence at least every 20ms. This is much worse outdoors due to the Doppler shifts and movement of users, in outdoor environments, where there is a lot of movement; the channel variability will increase to the extent that explicit beamforming will become irrelevant.

INTEROPERABILITY

As part of the training and sounding process in explicit beamforming, the AP and client must communicate and exchange information. To date, no interoperability between different vendors has been achieved. When we take into account the fact that AP chip vendors for Enterprise and Carrier products (Marvell, Qualcomm and Broadcom) are not the mass market client chip vendors (Intel, Ralink, Realtek) interoperability between vendors may be a long time coming.

INHERENT CONFLICT BETWEEN BEAMFORMING & MIMO

Both MIMO and beamforming are multi antenna technologies. However, there is an inherent conflict between the antenna requirements required by each technology to optimize performance.

Beamforming shape the beam direction by using constructive and destructive combinations of the same signal phase shifted from multiple antennas. The desired direction is achieved by changes to each antenna phase. In order for the signals to combine, they must have the same polarity.

The beamforming requirements are completely opposite from good MIMO requirements. In order to increase diversity in MIMO, we want signals from different antennas to have orthogonal polarity. You can optimize your antenna array for either beamforming performance or MIMO – but not both.

REGULATORY POWER

There are 2 types of beamforming from a regulatory point of view: narrow band beamforming, where the beams are created per tone/subcarrier and wide band beamforming, where a single beam is steered toward the client. In the 802.11 standards, only narrow band beamforming is standardized.

Because wide band beamforming focuses the energy in a single direction, they are considered to be 'greener'. To commend this, regulatory bodies, such as FCC in America, ARIB in Japan and MII in

China, have relaxed their requirements allowing wide band beamforming to transmit much higher EIRP power.

Unfortunately, narrow band beamforming cannot enjoy these regulations. By transmitting different beam per tone they spread the energy all over and cannot guarantee a narrow beam antenna over the entire bandwidth (different tones will have beams on different direction spreading the total beam all over).

THE HOLY GRAIL OF MULTI USER MIMO (MU-MIMO)

Across the industry, many people are waiting for the wave II 802.11ac chipsets, which will bring the long anticipated MU-MIMO functionality. MU-MIMO enables the AP to transmit data to multiple clients simultaneously using beamforming.

The need for this feature is obvious. High end APs with 4 antennas are capable of transmitting 4 spatial streams of information. However, on the receiving side there are simpler clients such as TV sets, storage, laptops, tablets and smart phones, which only have 1 to 2 antennas. Since the maximal number of streams is limited by the minimal number of antennas between the client and AP, the communication to 1 or 2 streams according to the simpler client. What a waste of perfectly good streams and potential bandwidth!

With MU-MIMO the AP will be able to fully utilize its capabilities.

For example, MU-MIMO is perfect for residential scenarios. Many houses have 2 to 4 TV sets and a similar number of PCs, tablets and storage devices. Most of these devices are static and don't move around the house very much. In particular, TV sets and storage, which consume high capacities, usually stay put in one place. This static scene with only a few clients is the ideal settings for MU-MIMO. Using MU-MIMO you can deliver 2 simultaneous video transmissions and double the link capacity.

How does MU-MIMO work?

MU-MIMO is achieved by forming a beam to each client. The beam has to be specifically formed to focus on the desired client and NULL the direction of the other client and all of its major reflections. In this manner, the transmitted information will be received well by the desired user. At the same time, the transmission will be significantly attenuated and, therefore, received in much lower power by the other MU users making it just a bit of noise. In this manner, all MU users enjoy good SNR (Signal to Noise Ratio) and get "almost" only the desired data stream.

Selecting fitting clients for MU-MIMO transmission and then creating the 2 beams and accompanying NULLS is not a trivial task. To address this, a new Scheduler block needs to be added to the AP. The Scheduler’s task is to identify the groups of clients out of all of the associated clients, which can share a link using MU-MIMO. To achieve this, much of the clients need to go through sounding, this information must be processed and only then can the scheduler create groups of possible candidates. Such candidates should be at a similar distance from the AP but in a different direction (angular spread) to achieve the optimal beam/null separation.

MU-MIMO is a true enhancement for residential Wi-Fi. However, for enterprise and carrier Wi-Fi networks it is virtually impossible to use. There are several reasons that contribute to this.

What they usually forget to you tell about MU-MIMO

SOUNDING OVERHEAD

To create possible matches for MU-MIMO, the scheduler needs repeated sounding information from all possible clients. This creates an enormous communication overhead. As a result, in slightly larger networks, the MU-MIMO overhead quickly exceeds any gains. Note that the schedulers work increases exponentially with the number of users, when trying to locate the suitable client groupings. In addition the sounding overhead in MU-MIMO is almost double in comparison to SU-MIMO since forming NULLS requires much higher accuracy than forming beams.

COHERENCE TIME & NULLS

Communication channels have a characteristic called coherence time, which identifies the required frequency of sounding feedback based on the channel's characteristics. In enterprise networks and, especially, in outdoor carrier environments the coherence time very low and clients have to repeat the sounding process every few milliseconds. In MU-MIMO the situation is even worse than TxBF while the time constants are much shorter. In the below MU-MIMO single and dual stream per client graphs you can see that after only 5ms the SNR is degraded to a level that requires sounding again. The graphs show the situation in a static indoor environment and will be much worse outdoors.

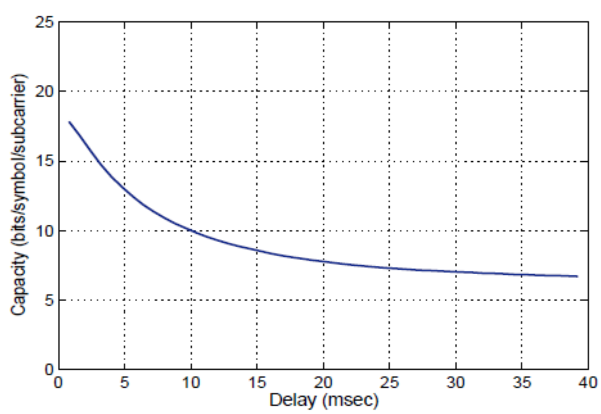

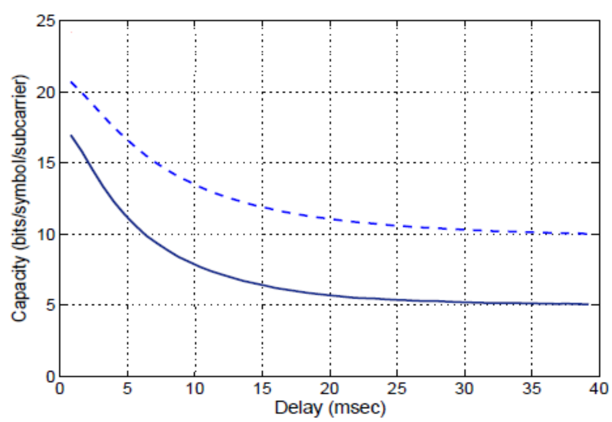

Over a short period (5-10 milliseconds) there is significant degradation in the signal. This is the result of the changing channel and the reason for frequent channel estimation updates using soundings. As you can see from the graph above, precision is critical to effective beamforming. In the case of NULLs it is even more so. NULLs unlike beams are steeper and very narrow. The NULL effect can provide a powerful more than 30 dB noise mitigation when you get it exactly right but almost nothing if the calibration is even slightly off by few degrees.

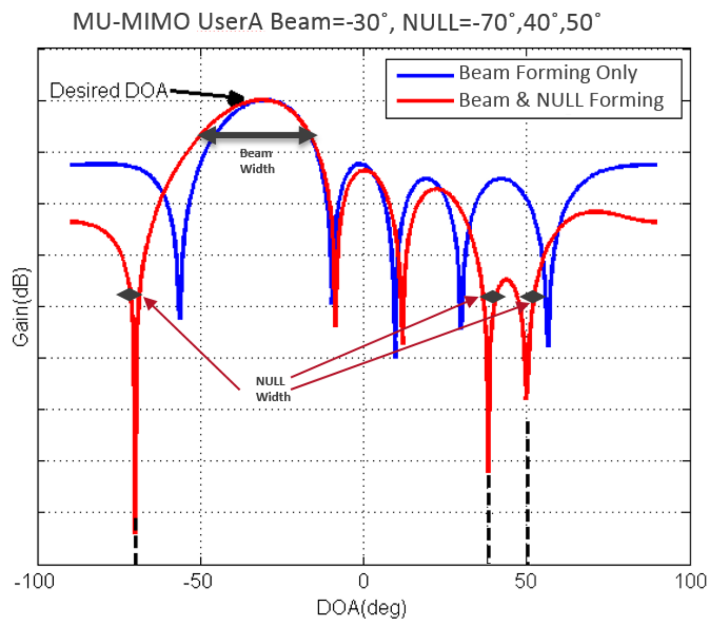

In the figure above we see a beam built by the TxBF system on one of the subcarriers toward a user called User-A in -30˚ direction (blue). In addition we see the same with added NULLs on direction -70˚, 40˚ and 50˚ representing reflections coming from User-B (red). We can clearly see the NULLs are much steeper than the beam enabling to get good enough separation but are much narrower in such a way that moving only few degrees to any side makes them ineffective.

DOPPLER

To exacerbate matters, in carrier networks users, vehicles etc. are constantly moving. This introduces a Doppler shift, a phenomenon that occurs when a signal 'hits' a moving object. To accommodate for these fast channel changes, the frequency of sounding exchanges becomes very high. At this rate, the MU-MIMO feature is rendered useless.

SNR & PERFORMANCE DEGRADATION

The performance experienced by MU-MIMO clients will be lower than if they were served alone. There are 2 major reasons for this.

First, the AP splits its output power between the served clients, which lowers the S in the SNR (each client receives a separate beam splitting the AP's energy into 2 beams with half the power per beam). Second, there is no perfectly quiet NULL. This means that each client experiences additional noise due to the other client's communication, which increases the N in the SNR. This degraded SNR can be so bad that it will make the feature impractical. For example, sending 450Mbps of traffic to single user without any overhead can be better than sending 130Mbps each with a lot of overhead to 2 users simultaneously.

THE MIMO MU-MIMO CONFLICT

When optimizing antennas for MIMO operation, our goal is to create as little correlation between antennas as possible. To achieve this we change the antenna polarity and separate them as far as possible. MU-MIMO, like beamforming, requires the opposite but even more so. To effectively create beams and NULLs, correlated antennas are much preferable. Since cannot optimize for both, compromises that lower performance have to be made for either or both SU-MIMO and MU-MIMO.

IS THERE A BETTER WAY?

Is there a solution that can overcome the limitations of the 802.11ac standard and deliver more of the promised speed and capacity?

To deliver high performance 802.11ac, we need to find a way to overcome the sounding overhead and still deliver focused beams (and NULLs) in the client's direction. To solve the challenge, let's break-down the tasks.

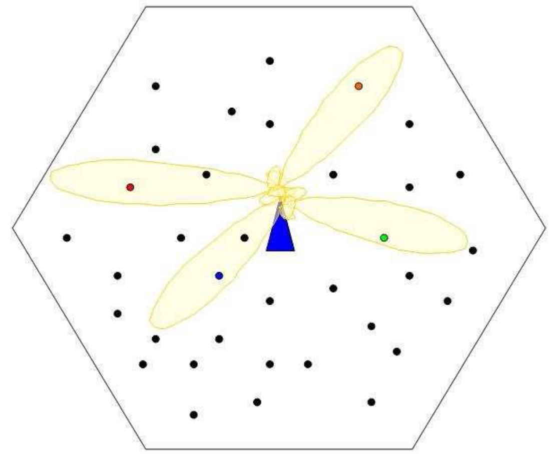

Let's assume that we know – through a secret method – the direction and distance of every user from the AP. And that this information is continuously updated. This means that we can easily select pairs of users that fit the requirements of MU-MIMO. Because the information is continuously updated, we can apply MU-MIMO successfully without the trial and error process. Such solution for MU-MIMO problem in dense user's environment is sometimes referred to as opportunistic beamforming with spatial structure.

Now, all we need to do is find the secret method of knowing the client directions and distances.

The answer to this is hidden in plain sight. In our initial discussion about beamforming, we mentioned implicit beamforming, where the AP is self-calibrating. This means that the AP is capable of measuring the channel for each client without requiring any overhead communication with the client. We also said that due to the complexity, the 802.11n implicit beamforming option was never used in standard off the shelf products. There is, however, a field-proven commercial implicit beamforming solution applying all the above.

For the purpose of solving most issues above, an ideal solution would be:

1) AP with self-calibrated array for no overhead implicit beamforming
2) Several such arrays working simultaneously in orthogonal polarities to achieve good MIMO polarity diversity without degrading beamforming performance
3) Spatial separation between such arrays adding space diversity on top maximizing MIMO capacity.
4) Scheduler capable of processing the implicit beamformer database for
    a. Reducing overhead and enabling MU-MIMO with any number of users
    b. Increasing beam/nulls performance and thus increasing MU-MIMO SNR
5) Wide Band beamformer with true narrow beam enabling the advantage of extra EIRP in regulatory area permitting such (ex. FCC, ARIB Japan, MII China and more)

GONET SYSTEMS XRF™ BEAMFORMING

GoNet Systems, a pioneer of Wi-Fi beamforming, has implemented a patent-based implicit beamforming technology called xRF™ Beamforming. The xRF Beamforming technology is fully interoperable with any standards based 802.11 b\g\n and now 802.11ac Wi-Fi client on both uplink and downlink. It delivers the benefits of beamforming – increased range, capacity and immunity to noise – in challenging indoor and outdoor environments.

Implicit beamforming – No client support required – No Overhead

One of the main advantages of implicit beamforming in comparison to explicit beamforming is client-side support. In explicit beamforming the AP and client must be fully interoperable. Implicit beamforming is transparent to the client and, therefore, requires no support on the client side for interoperability. With implicit beamforming there is no overhead. The AP can measure channel estimation information directly from the received packets and always stay calibrated in this way.

GoNet Systems employs wide band implicit beamforming. This means that the APs are allowed by many regulatory bodies to transmit at a higher power. The additional power compensates for the SNR losses enabling the APs to maintain high constellations and Wi-Fi rates.

The GoNet Systems AP has updated client information without requiring sounding for channel estimations. As a result it maintains high beam performance even in dynamic environments with frequent channel variations, where standard APs are barely functional. While standard beamforming updates its channel information using sounding every few milliseconds, GoNet Systems APs updates its database with each received packet. This approach enables the solution to easily support any dynamic, fast changing channel.

The updated client information is also used by the GoNet Systems AP as the input to the MU-MIMO scheduler. Without overhead levels that increase exponentially with the number of clients, GoNet System can perform MU-MIMO with up to 250 clients – even in tough outdoor environments.

ABOUT GONET SYSTEMS

GoNet Systems provides beamforming Wi-Fi solutions for 3G/4G data offload, Wi-Fi access and enterprise networks. With Next Generation xRF Beamforming technology, GoNet Systems GoBeam access points deliver x2-x4 the capacity and 90% effective noise mitigation in both indoor and outdoor deployments. The dual radio beamforming architecture that supports both 2.4 GHz and 5 GHz bands combined with a directional, automatic channel selection mechanism enables operators to optimize the frequency utilization. With a flexible architecture GoBeam platforms are designed to support the latest Wi-Fi standards 802.11n and 802.11ac.

GoNet Systems headquarters are located in Israel and the company has offices in Brazil, Japan, Korea and China. For more information visit our website at www.gonetworks.com.

GONET SYSTEMS