Welcome to an Understanding IEEE 802.11ad Physical Layer and Measurement Challenges

Note to presenter:
Edit this slide with your personal detail as shown

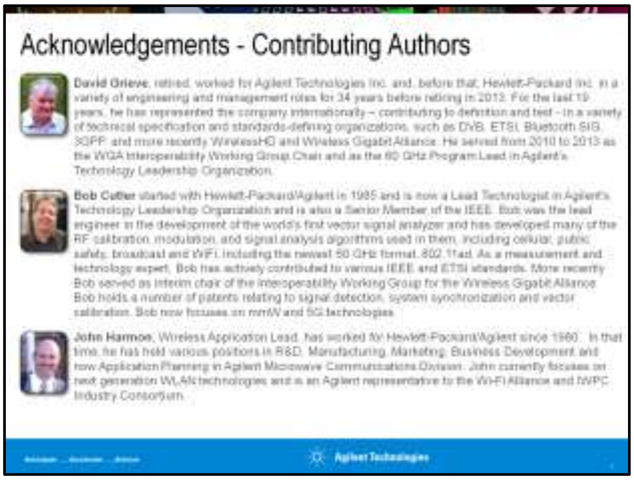

First, we should acknowledge the authors of this presentation for their extensive knowledge of wireless communications including 802.11ad and vast measurement expertise from Agilent's Electronic Measurements Group.

David Grieve, now retired, was our 60 GHz Program Lead and served as the Interoperability Working Group Chair in the WiGig Alliance.

Bob Cutler, a Lead Technologist in EMG's Technology Leadership Organization, developed all of the modulation analysis measurements and is one of the original inventors of the vector signal analyzer;

and John Harmon, a wireless application lead in the Microwave Communications Division.

These three represent 97 years of industry experience.

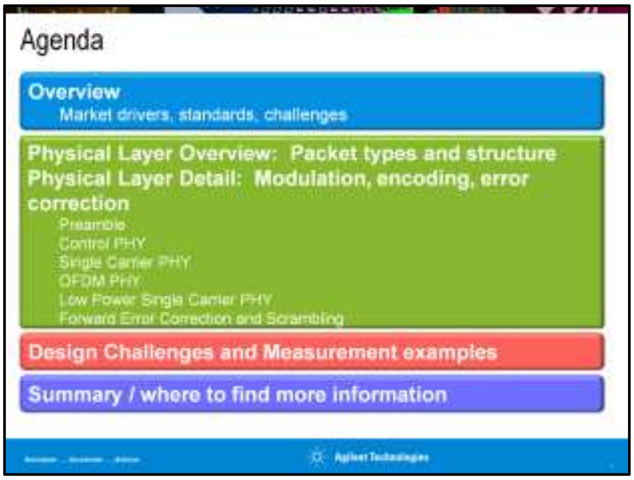

Our agenda today will start out with an overview to build some context on this WLAN standard

Moving right into the PHY packet structure, then we will dive into the details of the actual PHY packet

Along the way, we'll sprinkle in some real measurement examples to both illustrate concepts and to visualize what has been presented.

We'll finish up with some measurement examples and some concluding remarks

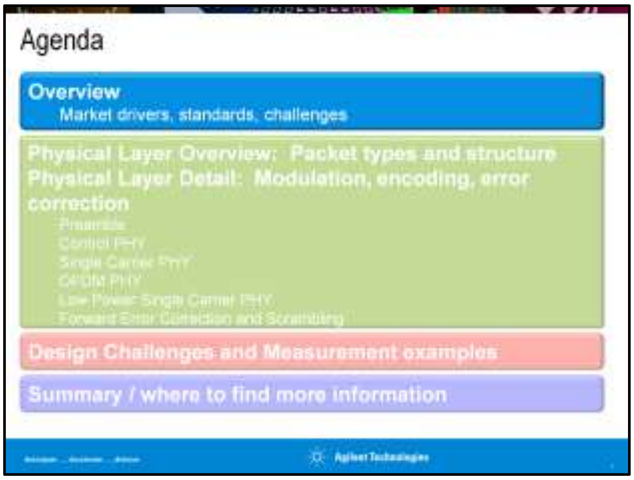

So let's get started…

By now everyone has heard of or experienced the tsunami of data in our world. Everything must be connected all people at all times.

This vision drives the WLAN market, which includes, but is not exclusive to, 802.11ad.

While WLAN is being integrated into more consumer products, there is demand for higher rates and simultaneous connections.

Enterprise adoption and Carrier-grade Wi-Fi raises the bar on system security and quality.

Wireless displays, multiple video streaming, rapid file upload/download, and transfer or streaming of large files all require significantly higher data rates.

And the Internet of Things demands more capacity.

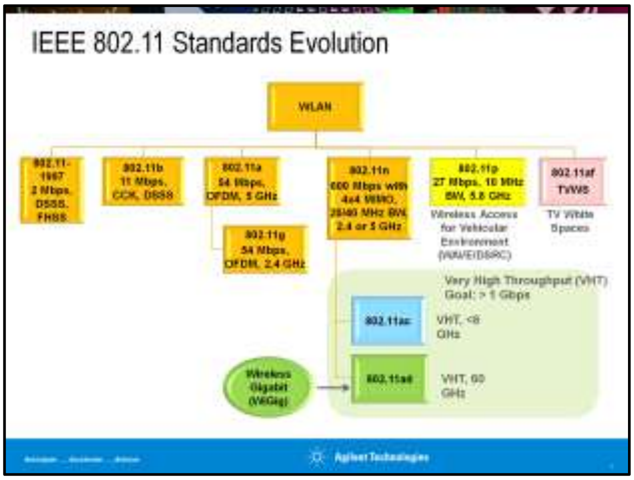

Since its initial introduction in 1997, several amendments to the IEEE 802.11 standard for different physical layers have been introduced.

This chart highlights the main amendments.

This presentation series is focused on 802.11ad specifically, so naturally we must also focus on its core mission, delivery of Very High Throughput rates.

The two IEEE project groups designated to provide “Very High Throughput” (VHT) are 802.11ac in the 5 GHz Band and 802.11ad in the 60 GHz Band. Both have aspirations of a minimum 1 Gbps throughput with top throughput of 6.9 (for 11ac) and 6.75 (for 11ad).

The goal is for all the 802.11 series of standards to be backward compatible, and for 802.11ac and ad to be compatible at the Medium Access Control (MAC) or Data Link layer, and differ only in physical layer characteristics

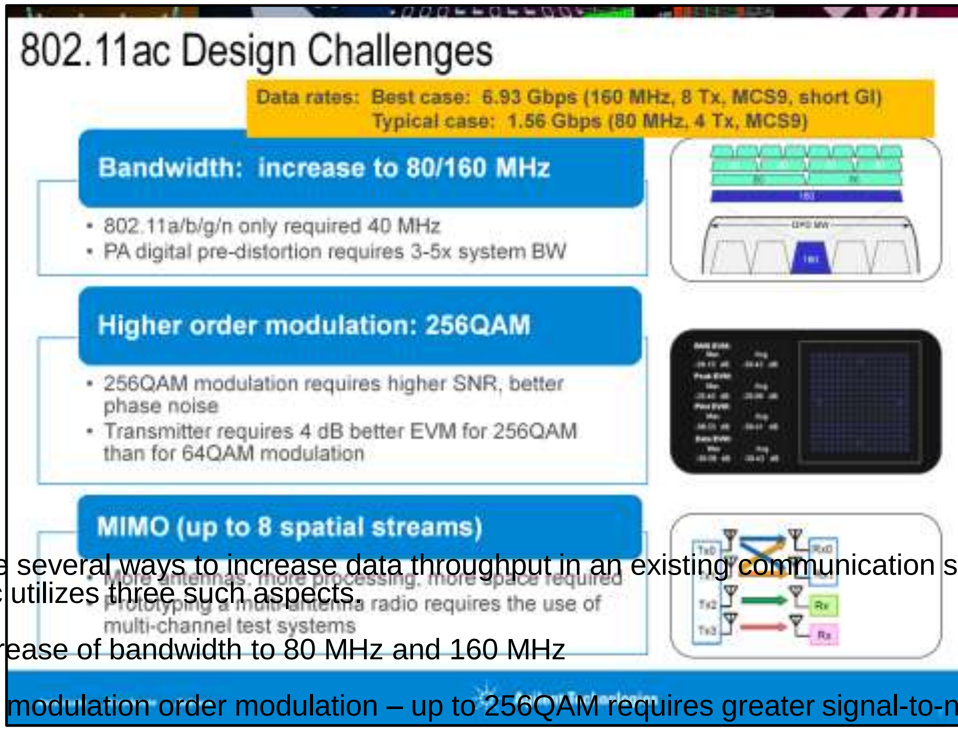

There are several ways to increase data throughput in an existing communication system. 802.11ac utilizes three such aspects.

- an increase of bandwidth to 80 MHz and 160 MHz
- higher modulation order modulation – up to 256QAM requires greater signal-to-noise ratio.
- Up to 8 potential spatial streams for MIMO.

All tolled, the system parameters combine to provide theoretical data throughput up to 6.9 Gbps – a remarkable speed for wireless.
<click>
However, we must put on our practical hats and lower our expectations to what will most likely be much less.

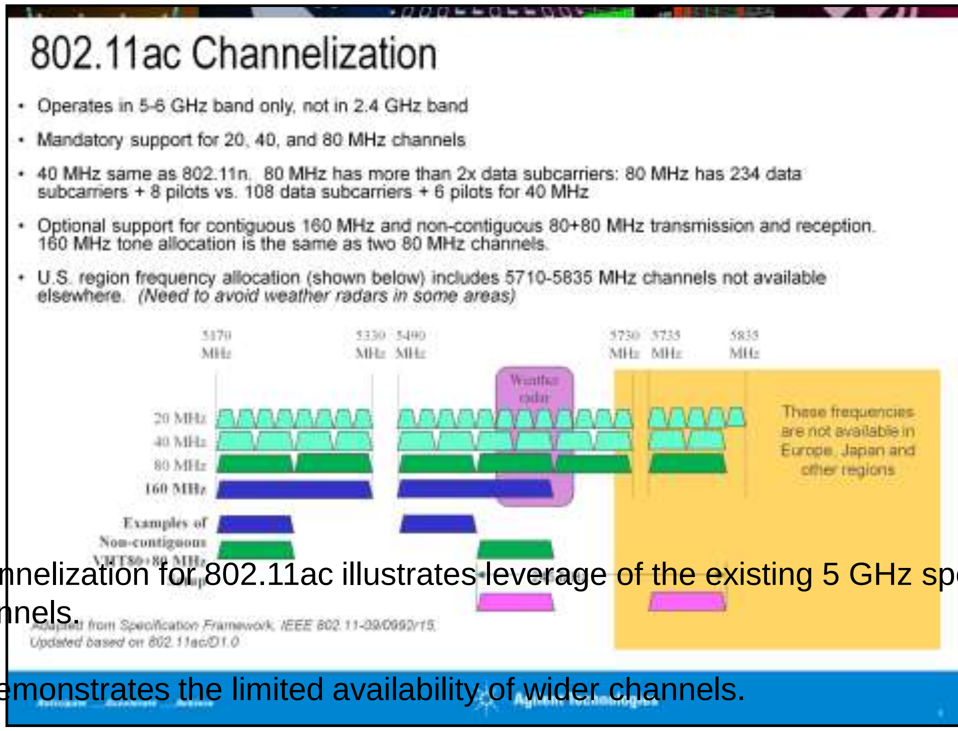

The channelization for 802.11ac illustrates leverage of the existing 5 GHz spectrum and channels.

It also demonstrates the limited availability of wider channels.

We can see that there are only two 160 MHz wide channels available, and the second channel cannot be used if there are weather radars nearby.

Thus the 80+80 MHz mode allows more opportunities for 160 MHz transmissions by combining two 80 MHz channels that are not adjacent in frequency.

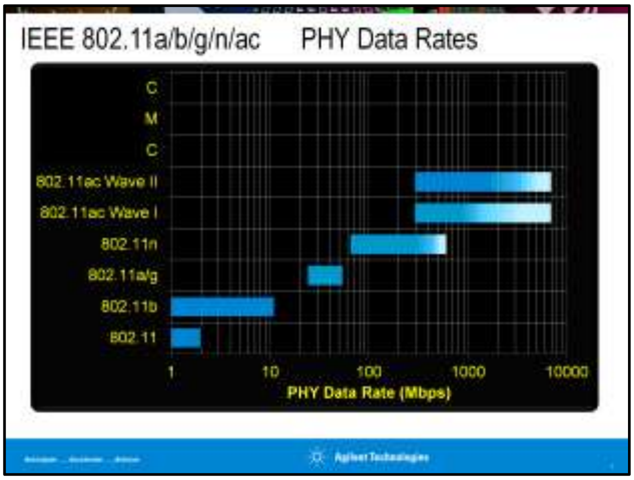

Nevertheless, the rates stated for 11ac vary from 300 Mbps to a top end of 6.9 Gbps, deployed in two waves.

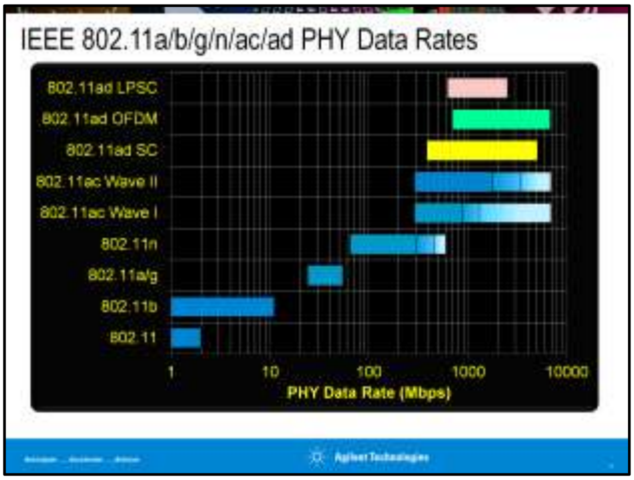

By comparison, 802.11ad will address higher throughput with three specific PHY using single carrier modulation, OFDM and a low-power mode.

We will see details of this later on.

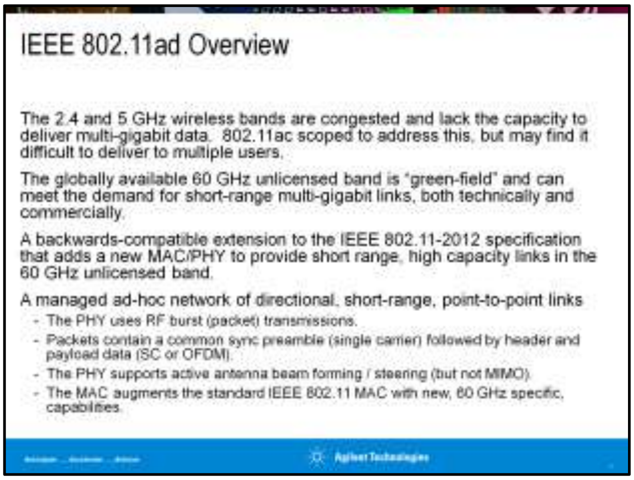

Because of the practical obstacles with lower band capacity, the IEEE looked to the unlicensed 60 GHz band to extend the 802.11 MAC and PHY definitions to support short-range (1m - 10m) wireless data interchange over an ad-hoc network at data rates up to 6.75 Gbps in the 60GHz unlicensed band.

It also supports session switching between the 2.4GHz, 5GHz and 60GHz bands.

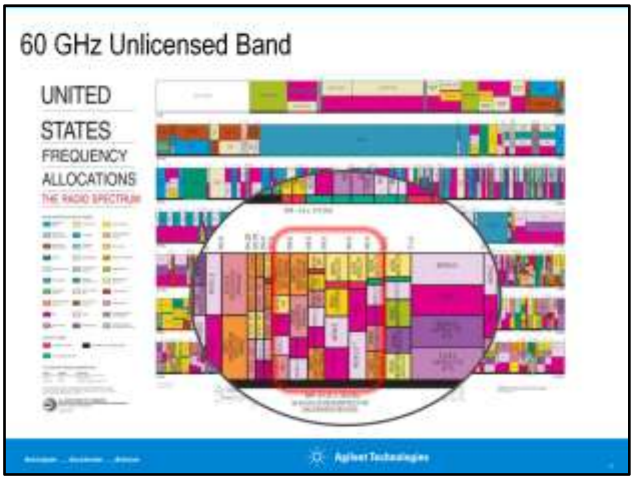

Since 1994, when the Federal Communications Commission (FCC) first proposed to establish an unlicensed band at 59-64 GHz, radio regulatory organizations around the World have been legislating appropriate frequency allocations and modulation parameters to enable similar unlicensed bands in their respective jurisdictions.

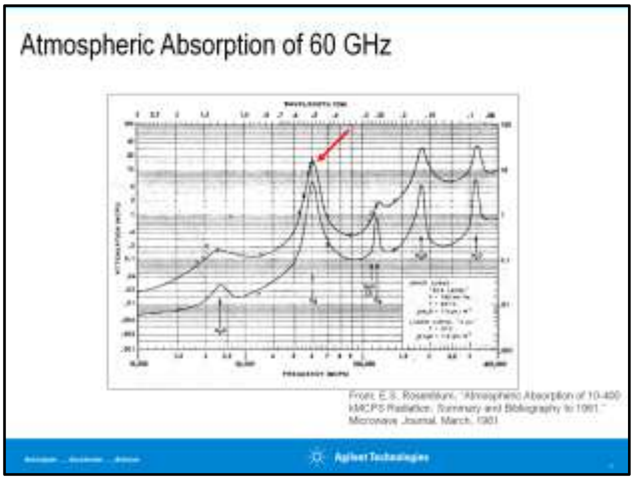

In the early 1960s, information was published about the unique absorption characteristics of higher frequency electromagnetic signals.

60 GHz is the absorption resonance for oxygen.

This is a particularly attractive property for limiting propagation in secure applications, but it is also useful in reducing co-channel interference.

Both of these properties are put to use for 60 GHz commercial use.

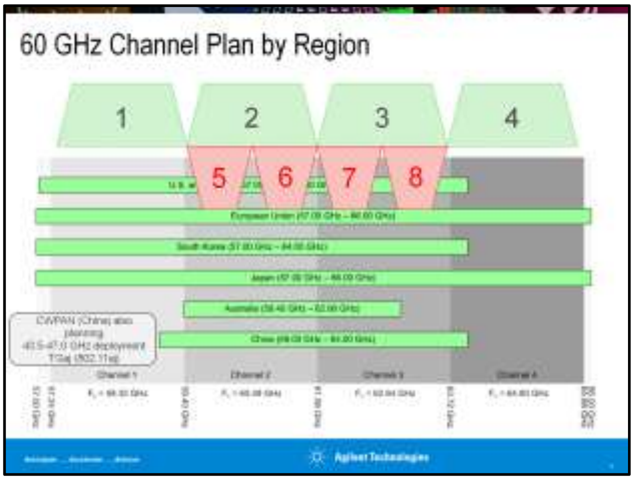

A channelization of the 57-66 GHz band has emerged by consensus from technical specification development work.

In November 2011 this channelization was approved by ITU-R for global standardization shown here.

Channel 2, which is globally available, is therefore the default channel for equipment operating in this frequency band.

As the diagram illustrates, not all channels are available in all countries.

<click>

China, for instance, is addressing this in two ways:

- By sub-dividing Channels 2 & 3 into 1 GHz channels 5 through 8
- And defining 802.11aj in the 43.5-47 GHz band with compatible 11ad MAC and PHY.

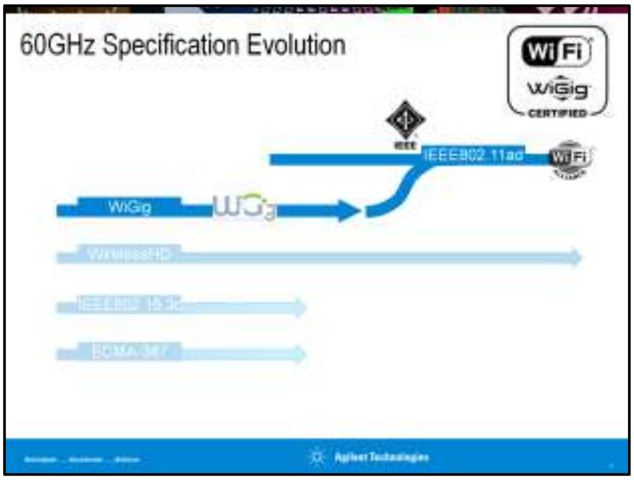

WirelessHD has the distinction of being the first technology to commercialize the 60 GHz space and is still commercially active today.

IEEE 802.15.3c and ECMA-387 made attempts to define standards in the 60 GHz band, but neither has seen any commercial development.

<click>

WiGig was the latest standardization initiative.
Since its incorporation in 2009, the WiGig Alliance designed to address a wide range of use cases using both Single Carrier and OFDM PHYs, but implemented with a much higher level of commonality across the two modes and enjoys support from multiple vendors.

<click>

It has been adopted unchanged by IEEE on that basis for the WLAN 60 GHz MAC/PHY. This is a key point; the IEEE 802.11ad and WiGig MAC/PHY documents have been essentially identical from the start.

In December 2012, the Standards Board approved and released 802.11ad.

As such the Wi-Fi Alliance will now take the lead responsibility for certification and marketing of the standard.

<click>

802.11ad devices will be certified as WiGig Certified.

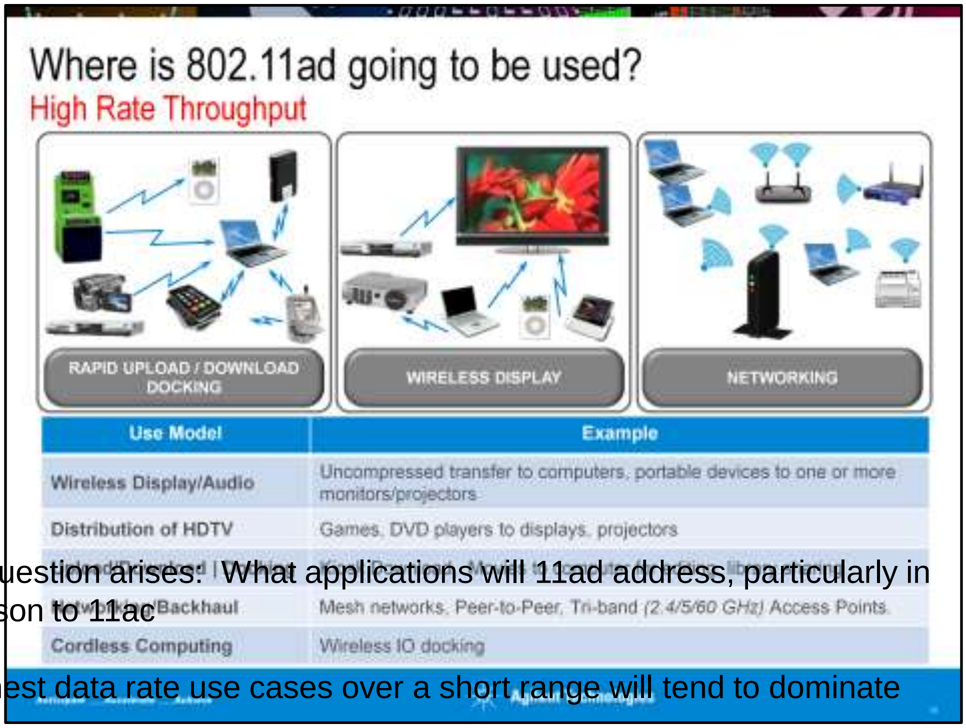

So the question arises: What applications will 11ad address, particularly in comparison to 11ac

The highest data rate use cases over a short range will tend to dominate

Upload/download, docking, Uncompressed and multiple streaming video rise to the top.

But networking in concert with the lower 2.4 and 5 GHz bands, docking and small cell backhaul
will also exist.

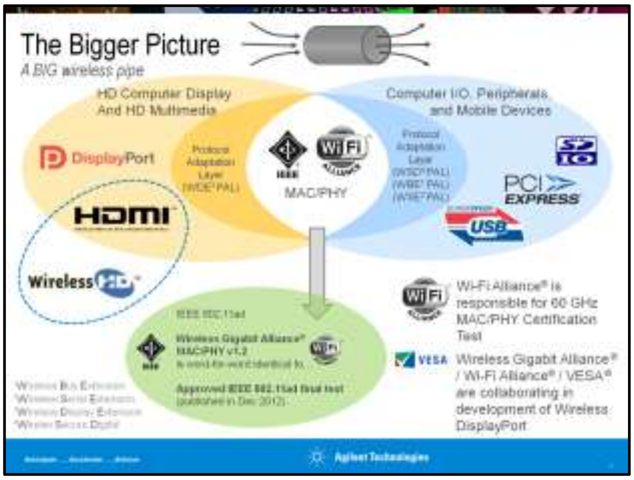

Wireless communication in its simplest form is nothing more than a replacement of what we currently enjoy in a wired form, thus liberating us from the bounds of the physical wires while remaining connected.

<click>

We can imagine this as a big wireless pipe. Something has to go into the pipe and something has to come out.

In the case of networking, it is a wireless implementation of gigabit Ethernet.

<click>

The situation is similar for other high-speed digital serial connections.

<click>

All of these ride as payload applications upon the PHY and MAC of the communication system.

Originally developed by the Wireless Gigabit Alliance privately, several Protocol Application Layers, or PALs, have been defined in concert with the alliance or consortium associated with that particular standard to ensure these transitions are transparent.

<click>

The Wireless Display Extension PAL has been written to support both DisplayPort and HDMI connections

<click>

HDMI-branded cable replacement is based on WirelessHD technology.

<click>

The Wireless Bus Extension PAL has been written with PCIe in mind;

Wireless Serial Extension PAL has been written with USB 3.0 in mind.

It is important to emphasise that, while the groundwork has been done in WGA to facilitate wireless versions of these established brands, whether they actually happen will depend on successful negotiations with the owners of these brands.

<click>

A liaison has been established with VESA to facilitate the deployment of this technology as a wireless version of the “DisplayPort” brand.

Other liaison agreements are in process, but confidential.

<click>

The formal agreement by the Wireless Gigabit and Wi-Fi Alliances establishes that going forward the Wi-Fi

Alliance will facilitate the deployment of this technology under the “Wi-Fi” brand.

<click>

As such, The Wi-Fi Alliance will be the certification body for 11ad MAC and PHY.

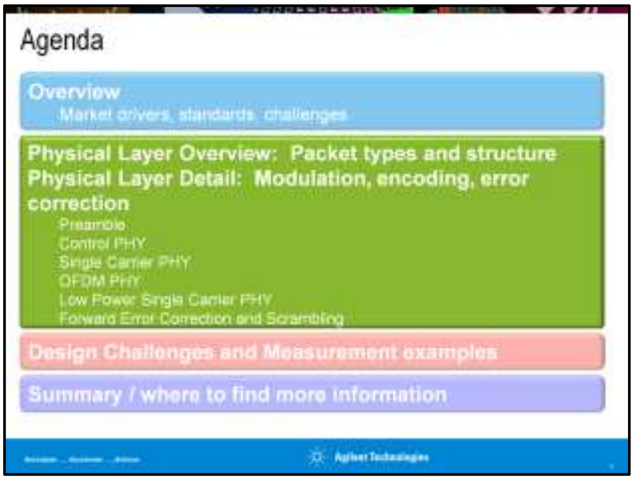

Let’s take a look at the Physical Layer

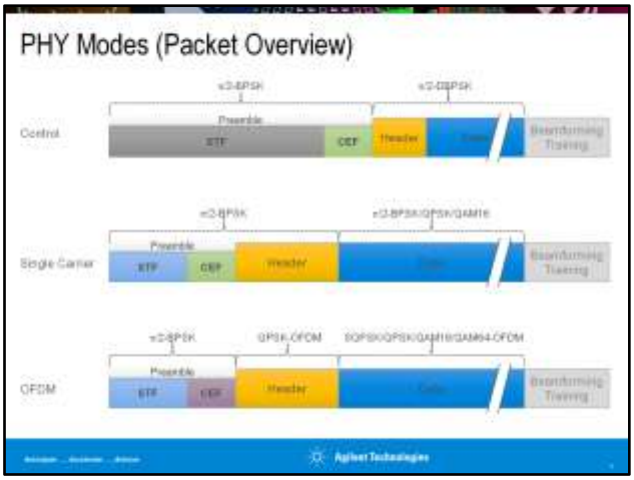

The PHY supports three distinct modulation methods:

1. Spread spectrum modulation; the very robust Control PHY.
2. Single carrier (SC) modulation; the Single Carrier PHY and the Low Power Single Carrier PHY.
3. Orthogonal Frequency Division Multiplex (OFDM) modulation; the OFDM PHY.

Each PHY type has a distinct purpose and packet structure but care has been taken to align the packet structures, and in particular the preambles, to simplify signal acquisition, processing and PHY type identification in the receiver.

The three PHY types share an essentially common preamble structure comprising a Short Training Field (STF) followed by a Channel Estimation Field (CEF). These fields are based on repeating Golay sequences and will be described in more detail shortly.

In all cases, the preamble is followed by a header field conveying information about the rest of the packet. Most importantly it indicates the specific Modulation and Coding Scheme (MCS) being used for the payload part of the packet.

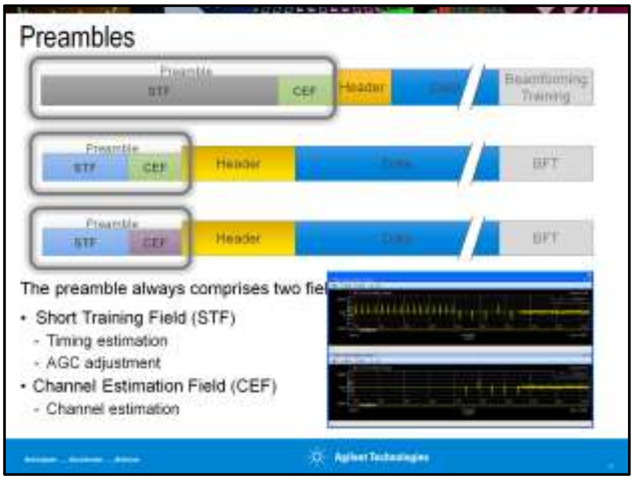

The Preamble has two roles:

- Short Training Field for timing and amplitude adjustment
- and the Channel Estimation Field

Note that the Control PHY Short Training Field is slightly longer than the other PHY. This is for robustness.

The graphic in the lower right is an output waveform of a correlator for Golay Complementary Sequences – in this case for length of 128.

Let’s take a closer look at Golay Sequences.

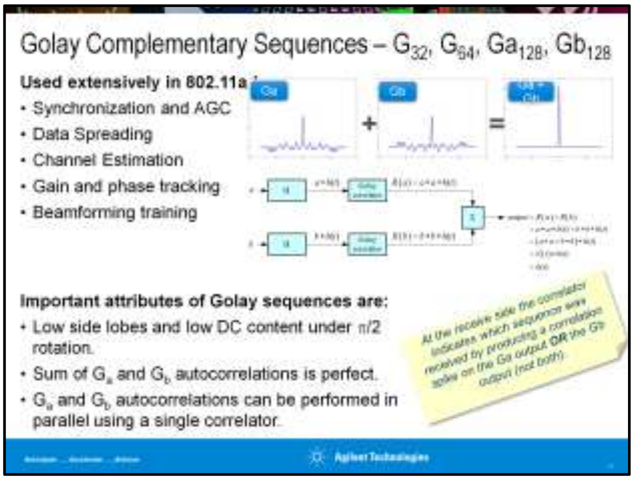

Golay complementary sequences are used extensively in 802.11ad and are something of a “miracle cure” in that they can provide a suitable signal for:

- clock synchronization and AGC adjustments;
- data spreading sequence to provide coding gain, processing to compute an estimate of the channel response.
- They are also used as a “guard interval” (GI) in the SCPHY, which is, in effect, a “pilot” gain and phase tracking symbol
- and in beamforming training.

Golay codes are sequences of bipolar symbols (±1) that have been mathematically constructed to have very specific autocorrelation properties.
The 'a' and 'b' suffices indicate that the sequences form a complementary pair and the numerical suffix indicates the sequence length.
802.11ad uses Ga_{32}, Ga_{64}, Ga_{128} and Gb_{128} sequences.

The mathematics of sequence construction is relatively simple, but beyond the scope of this presentation. It is sufficient to know that three important attributes of these sequences are,

1. The autocorrelation of each sequence has very low side lobes (false correlation peaks) and low DC content under $\pi/2$ rotation
2. The sum of the very good (but imperfect) autocorrelation functions of the Ga and Gb sequences is perfect - the side lobes cancel exactly
3. The Ga and Gb autocorrelations can be performed in parallel using a single, hardware efficient (and therefore fast) correlator.

Background notes – not used in narrative

At the transmit side, we transmit **either** the Ga sequence **or** the Gb sequence as required (they are not transmitted at the same time). At the receive side the correlator indicates which sequence was received by producing a correlation spike on the Ga output **or** the Gb output (not both).

The basic principle behind the structure of the channel estimation fields is illustrated here.

We have two time sequences, *a* and *b*. If we pass sequence *a* through the channel **H**, we convolve the sequence and the channel impulse response, *h(t)*.

If we pass the received signal through a Golay correlator for the known input sequence then we get the autocorrelation of sequence *a* convolved with the channel impulse response.

If sequence *b* is processed similarly, we get the autocorrelation of sequence *b* convolved with the same channel impulse response.

If we add the two results together then, because sequences *a* and *b* are Golay complementary sequences, the sum of their autocorrelations is an impulse response and we are left with the channel response, *h(t)*.

Here, we see the structure of the three different preamble types in more detail, illustrating that in each case the basic building blocks are the Golay complementary sequences Ga_{128} and Gb_{128}.

The short training field provides a signal for clock synchronization and AGC adjustment.

These tasks only require the regular, nominally constant, amplitude “tick” of the correlation peak from a single correlator, so in this case 802.11ad uses the complementary nature of the Ga_{128} and Gb_{128} sequences to also signal packet type information.

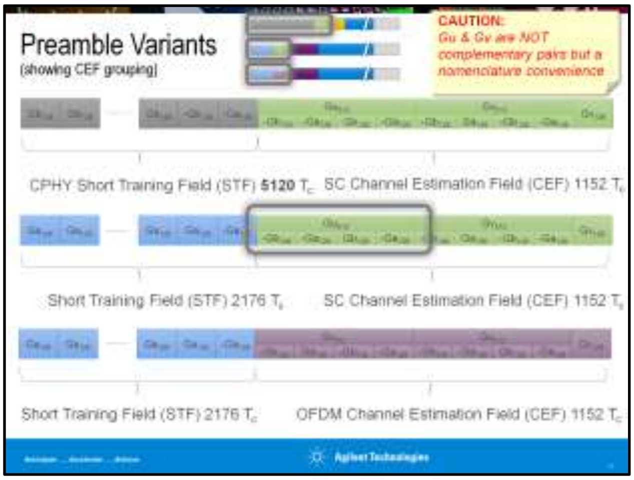

The construction of the CEF is more mathematically sophisticated than the STF and does take advantage of the Golay sequences' complementary property; the first four and second four Ga_{128}/Gb_{128} complementary sequences are logically grouped into 512-symbol field groups labeled Gu_{512} and Gv_{512} .

<click>

Each group provides a concatenation of two 256-symbol complementary sequences which can be used for channel estimation (as will be explained shortly). So the Gu_{512} and Gv_{512} groups can be used to perform two independent channel estimations. If necessary, these estimates can also be summed to give a composite result.

<click>

It's important to realize that the Gu_{512} and Gv_{512} labeling is a slight corruption of the pairing nomenclature. The Gu_{512} and Gv_{512} groups are not themselves a complementary pair of sequences, they are, however, 512 symbols long as the suffix suggests.

The main purpose of this grouping and labelling is to illustrate the fact that there are two channel estimation opportunities in the CEF and also that the field ordering is used to signal the payload modulation type. In an SCPHY packet the order is Gu_{512} then Gv_{512}, while for an OFDMPHY packet the order is reversed.

This time ordering of Gu_{512} and Gv_{512} is unimportant for channel estimation since each produces an independent estimate.

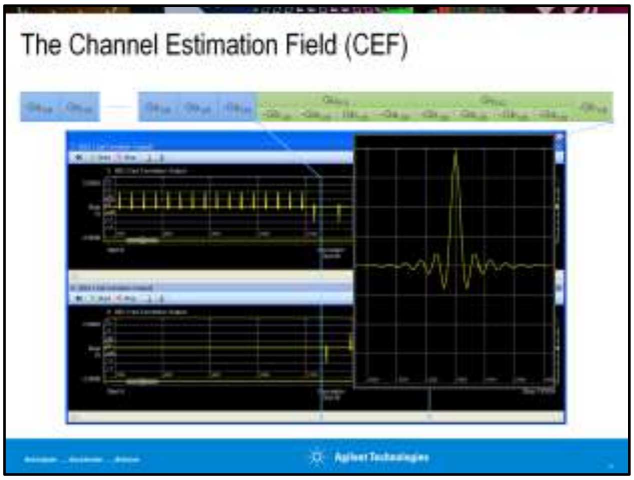

This slide shows traces of the Ga_{128} and Gb_{128} correlator outputs for a very clean 802.11ad SCPHY packet.

As we will see later in this presentation, a great deal of information can be gleaned from examining the less than perfect correlator outputs from a practical signal.

A minor, but important observation, is that the correlator output is not a simple impulse spike, but more of a SinX/X waveform, the significance of which will become obvious when we look into correlator measurements later on.

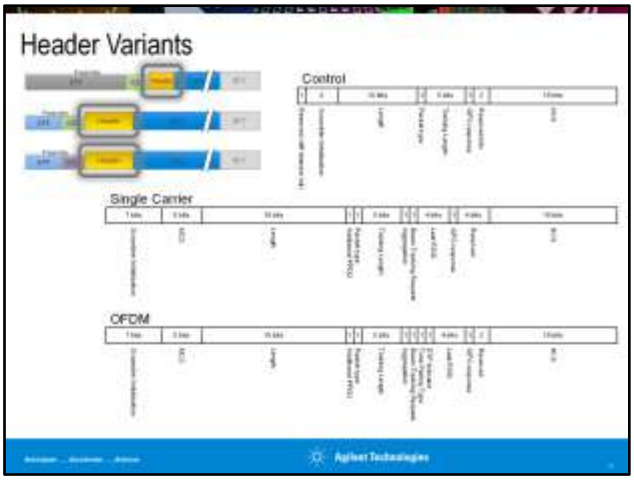

The Header Fields convey the necessary information to the receiver to enable proper capture, demodulation and decoding of the data field.

We will not go into each of these fields here.

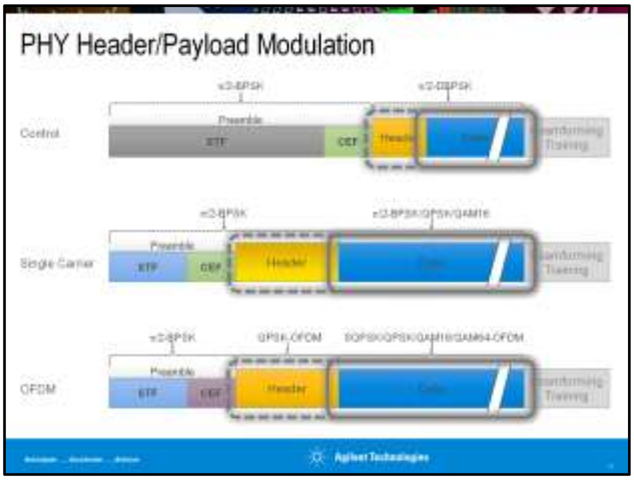

Now let’s examine the Payload

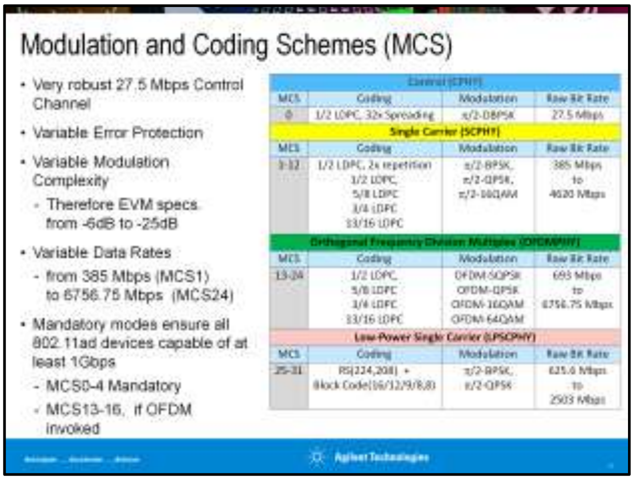

802.11ad tabulates 32 different modulation and coding schemes. However, as we have seen, there are just a few variations in the modulation and encoding of the preamble and header fields across all 32 MCS.

Thus we can quickly simplify the picture by dividing the MCS list into four basic classifications as shown here.

The control PHY emphasizes reliability over raw speed in order to establish a link.

Single Carrier and Low Power Single Carrier packets use the same modulation and so the measuring receiver can demodulate “blind” then decide which was sent based on the data encoding. In a DUT, the use of MCS 25 – 31 may be a negotiated decision.

It is clearly important, for the reliable establishment and maintenance of connectivity, that the control channel should be as robust as possible.

It is perhaps less clear why so many MCS are required.

Within each of the SC, LPSC and OFDM categories, the specific MCS selects a different pairing of error protection coding and modulation depth, which taken together provide the user with a logical progression of link quality versus throughput operating points.

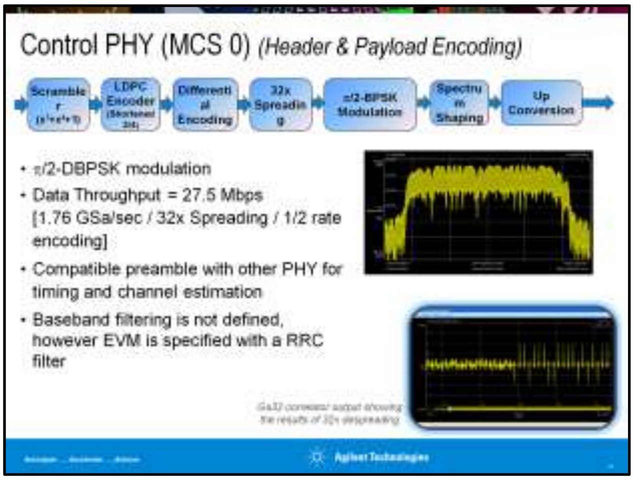

Modulation and Coding Scheme 0 (MCS0) is by far the most robustly coded (and consequently, lowest throughput) mode.

Its purpose is exclusively to transmit control channel messages and it is referred to as the Control PHY (CPHY).

Support for MCS0 is mandatory in 802.11ad.

The CPHY robustness is evident from its use of differential encoding, code spreading and BPSK modulation.

Differential encoding eliminates the need for carrier tracking, 32x spreading contributes a theoretical 15dB gain to the link budget, and BPSK is, of course, very noise tolerant.

A Ga_{32} Golay complementary code is used as the spreading code, so we can see the result of the despreading process directly by looking at the Ga_{32} correlator output as shown in the inset graphic.

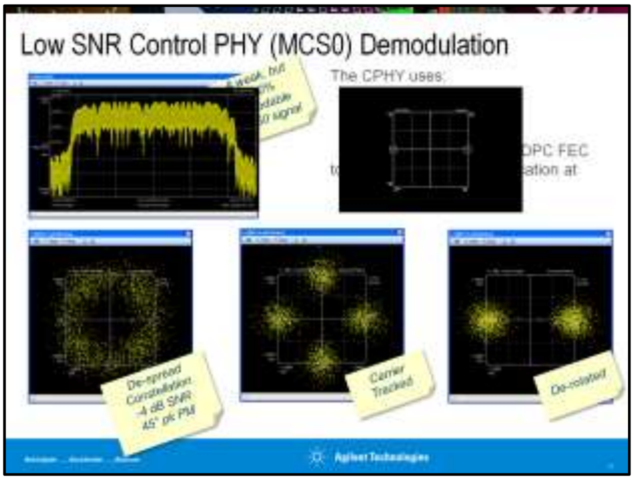

The 802.11ad CPHY is remarkably robust. As we described earlier, it uses differential encoding, code spreading, and BPSK modulation, augmented in the decoding step by a rate 1/2 LDPC FEC to ensure reliable communication at very high path loss.

Here we see a 30 dB signal-to-noise ratio Control PHY signal.

<click>

If we drop the signal to noise, we can see in the adjoining I-Q display, the recovered constellation gets noisier.

<click>

<click>

Here we have a negative 4 (-4) dB SNR. As you can see, the signal is for practical purposes “invisible”, its presence might be revealed by lengthy averaging on a carefully adjusted spectrum analyzer, but it would certainly be undetectable using a broadband power meter, yet this CPHY signal can be successfully decoded with a header checksum pass.

<click>

The first constellation shows the post despreading constellation, there may be a signal there but it appears to be suffering from significant phase jitter. In fact, this example signal was generated with 45° pk phase modulation at 1 MHz (as well as the -4 dB SNR). However, because MCS0 uses differential encoding it is very tolerant of phase jitter. To give a clearer picture of the constellation from a state decision point of view, we can turn on carrier tracking. A commercial receiver will probably not implement carrier tracking, but for a measuring receiver this is a useful diagnostic for separating out and understanding the nature of the transmitter’s impairments.

<click>

The second constellation shows the de-spread constellation with carrier tracking enabled. At first glance it appears to be QPSK, but of course it is pi/2 rotating BPSK. By default we would look at the rotated constellation because this can reveal any quadrature error or IQ distortions by keeping the I and Q components of the modulation distinct. However, it can also be useful to see the modulation as a BPSK constellation to subjectively assess the state decision noise margin.

<click>

The third constellation shows the de-rotated, carrier tracked, de-spread constellation of the signal whose spectrum is illustrated in the top graphic.

<click>

It is now perhaps easier to believe that such a weak signal can be successfully decoded error free. Even at -4dB SNR, the constellation is clean enough that few, if any bit errors are being corrected by the LDPC decoder.

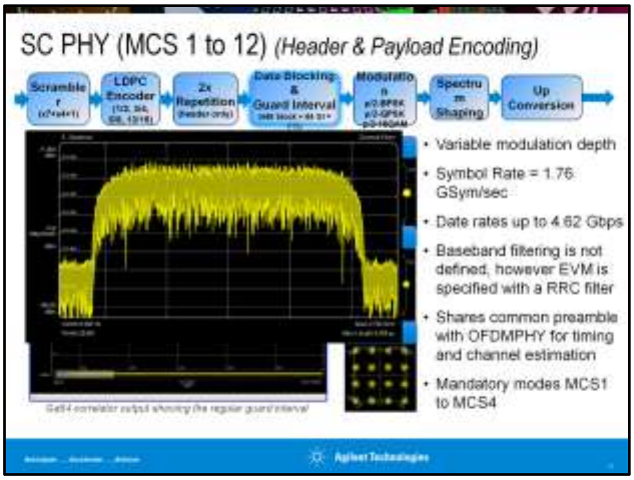

Modulation and Coding Schemes 1 through 12 (MCS1 – MCS12) employ single-carrier modulation; specifically BPSK, QPSK or 16-QAM modulation of a (suppressed) carrier at the channel center frequency, at a fixed symbol rate of 1.76 Gsym/s.

<click>

All 12 modes are essentially identical in their channel encoding steps.

They differ only in the choice of error protection ratio and modulation density to tradeoff between throughput and robustness for up to 4.62 Gbps throughput.

Support for modes MCS1 to MCS4 is mandatory.

<click>

An important step in the encoding process is the data blocking and guard interval.

<click>

This divides the modulation symbols into groups of 448 symbols interspersed with 64 symbol "guard intervals" (GI) that provide the receiver with a periodic known reference signal to assist with gain and phase tracking.

The 64 symbol guard interval is a Ga_{64} Golay sequence and its periodic occurrence can be confirmed by examining the output of the Ga_{64} correlator.

<click>

Given the symbol rate, and cost constraints, the use of the Guard Interval provides the option of a low-implementation-cost carrier tracking mechanism.

The modulation is very conventional single-carrier modulation which is $\pi/2$ rotated to minimize the peak to average power ratio (PAPR) of the BPSK modulation (the GI's are always BPSK modulated) and to allow equivalent GMSK modulation.

Spectrum shaping is mandated but the details are not specified, to permit some design freedom.

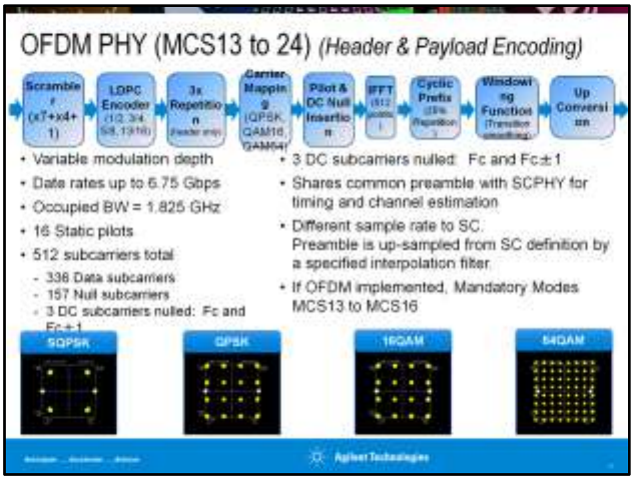

Modulation and Coding Schemes 13 through 24 (MCS13 – MCS24) employ multi-carrier modulation; specifically Orthogonal Frequency Division Multiplex (OFDM) modulation, which can provide higher modulation densities and hence higher data throughput than the single carrier modes at the expense of hardware complexity, power consumption, linearity and phase-noise.

The OFDM modulation is straightforward.

- It is based on a 512-point FFT in which 336 carriers are allocated to data.
- There are also 16 fixed pilots distributed across the spectrum
- The central 3 tones and Fc and Fc +/-1 are suppressed to avoid any issued with DC leakage.
- The remaining 157 tones are nulled.
- There is only one cyclic prefix length (25%).

As already mentioned, the OFDM and SC PHYs share a common preamble, but their reference sampling rates are different.

As for the single carrier modes, all 12 OFDM modes have near identical encoding, varying only in choice of error protection ratio and the depth of modulation applied to the OFDM data carriers.

Support for OFDM modulation is not required by the specification, but if it is implemented, then MCS13 to
MCS16 are mandatory to ensure some level of interoperability between

OFDM-capable devices.

The LDPC encoding is identical to that used in the single carrier modes.

<click>

Here, the individual OFDM carrier modulation may be SQPSK, QPSK, QAM16 or QAM64.

The QAM16 and QAM64 modulations are conventional, but a closer look at the SQPSK and QPSK constellations shown should raise some questions.

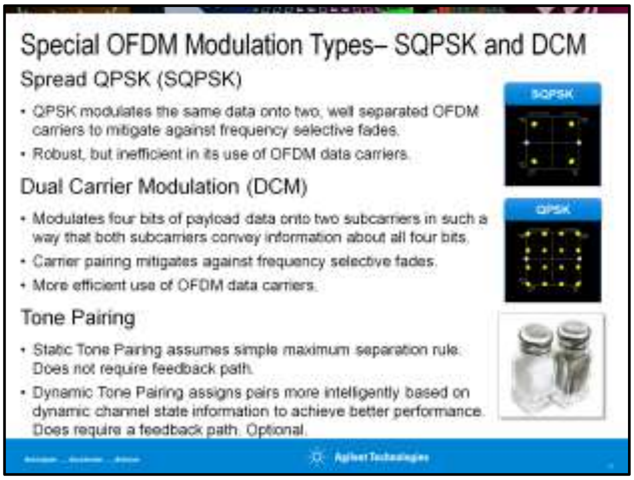

In Spread QPSK we QPSK modulate the same data onto two, well separated OFDM carriers. The objective is to mitigate against frequency selective fades.
This gives us a robust signal, but with limited throughput because of its inefficient use of OFDM data carriers.

<click>

In Dual Carrier Modulation we modulate four bits of payload data onto two subcarriers in such a way that both subcarriers convey information about all four bits.
This retains the carrier pairing advantage to combat frequency selective fades, but it also gives us the next level of throughput because of the more efficient use of OFDM data carriers.

<click>

Both SQPSK and DCM pair off carriers (or tones) to achieve frequency diversity.
11ad supports two methods of pairing tones.

Static Tone Pairing assumes a simple maximum separation rule. This will often, but not always, be a good, pragmatic choice. The advantage over dynamic pairing is that it does not require a feedback path from the receiver to the transmitter to communicate the current pairing, since it is fixed.

Dynamic Tone Pairing assigns pairs more intelligently based on dynamic channel state information to achieve better performance. It <u>does</u> require feedback from the receiver. Not all receivers will want to provide the necessary channel information so it is an optional feature that is negotiated between the transmitter and receiver.

Additional background:
SQPSK is Spread QPSK, in this mode, the OFDM carriers are paired and the same data is modulated onto two carriers maximally separated in frequency to improve the modulation's robustness in the presence of selective frequency fading. The idea is that if one carrier is lost to a null the other is unlikely to be affected at the same time. The pairing of tones is normally static, but there is an option to pair them dynamically according to channel conditions, which has been shown to provide additional robustness. In the SQPSK image, the yellow constellation dots are the QPSK data modulation points while the white dots identify the BPSK pilot modulation. The signal illustrated has 25 dB SNR.

The QPSK modulation employed in 802.11ad is actually Dual Carrier or DC-QPSK, which is a form of diversity coding. Data is taken 4 bits at a time to create two QPSK symbols which are then combined by a matrix operation to produce two Dual Carrier Modulation (DCM) symbols that, as with SQPSK, are mapped to widely separated subcarriers. Taking the possible modulation states of each carrier separately there are 16 valid phase/amplitude states in each DCM symbol and so the OFDM constellation looks cumulatively like 16-QAM, however, this method of modulation actually preserves the QPSK minimum distance between states and so has the same performance as QPSK in the presence of Gaussian noise, and has improved performance in the presence of fading due to the frequency diversity.

The yellow constellation dots shows the 16 possible locations for the DCM data modulation points while the white dots identify the BPSK pilot modulation. The signal shown in the (DC-)QPSK image has 25 dB SNR.

The yellow constellation dots are the QAM16/64 data modulation points while the white dots identify the BPSK pilot modulation. The signal shown in the QAM16 image has 30 dB SNR, while the signal shown in the QAM64 image has 35 dB SNR.

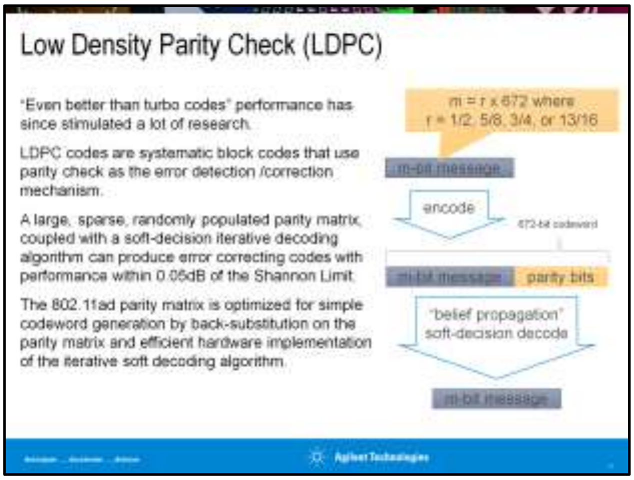

Low Density Parity Check encoding was first proposed by Gallager in his 1960 Doctoral dissertation, but was then largely ignored for 35 years! Recent research is now applying these to real communication systems, e.g. 802.11.

LDPC codes are systematic block codes that use parity check as the error detection /correction mechanism. Gallager and others found that a large, sparse, randomly populated parity matrix, coupled with a soft-decision, iterative-decoding algorithm could produce error correcting codes with performance within 0.05 dB of the Shannon Limit.

A potential problem with LDPC codes is that their performance only becomes asymptotically good with large data blocks. Fortunately at 60 GHz, large data blocks are "standard."

Although 11ad does not use MIMO, it does employ beamforming training as an addendum to the packet.

<click>

As discussed earlier, this field utilizes Golay sequences to convey this information – Ga64 and Ga & Gb 128.

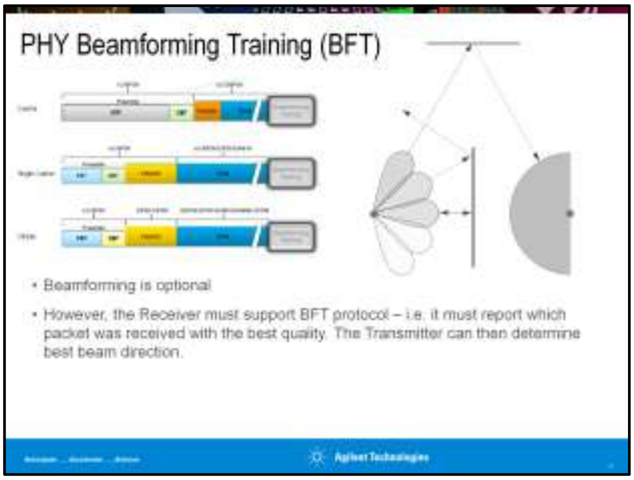

While beamforming is optional, all receivers must support the MAC level messaging to indicate which received packets were best, so that if beamforming is being used, the transmitting device will know which beam to select.

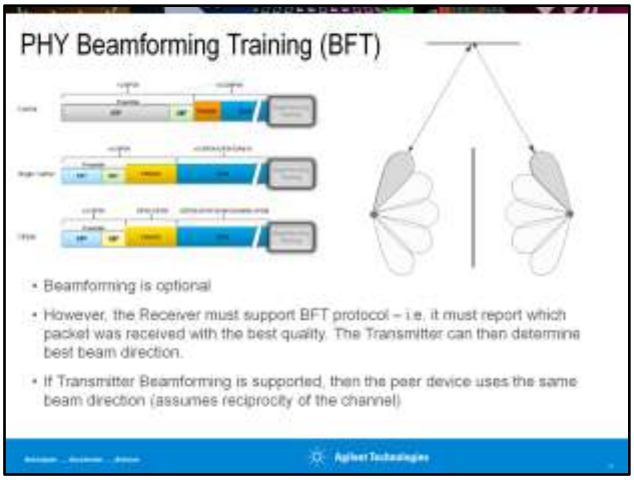

If transmitter beamforming is supported, a device can use feedback from its peer to determine which beam is best. If both devices are transmit- beamforming-enabled, then there is an assumption of reciprocity in the beam path selection.

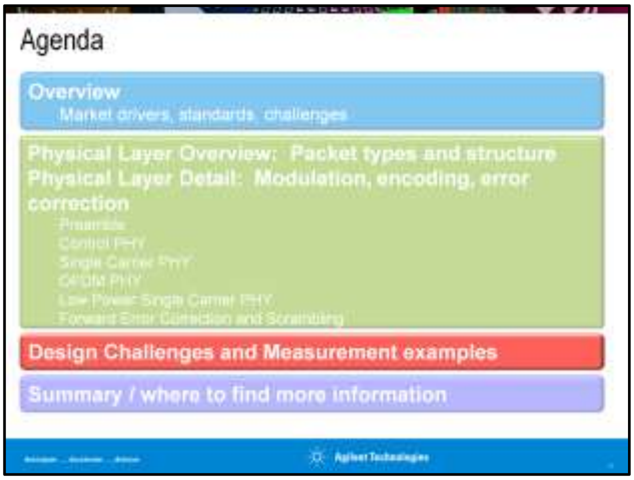

Let's take a closer look at measurement challenges and specific measurements

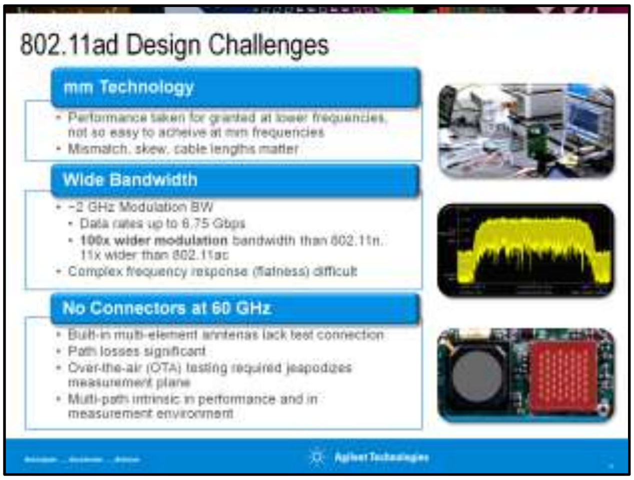

802.11ad has its design challenges as well.

First, it is millimeter. 20 years ago, we worried about 5 GHz technologies. There were no fast ADCs or DACs, mismatch and skew were problems, antenna design was a black art, computation was too slow. Now we can buy these chips very inexpensively.

Fast forward to today, 60 GHz, like most millimeter, presents the same challenges of 20 years ago. We have to use I-Q topologies of create an intermediate frequency. Antenna design at millimeter frequencies is not something anyone can do. Cable lengths at millimeter are critical.

But, it is also better. These circuits are done in CMOS, which is remarkable. SiGe is also used for higher power – mostly backhaul applications. And on-chip computation is very fast now.

Bandwidth is an issue. Wide BW circuits are not trivial and require flatness corrections.

Antennas, like this multi-element beamsteering antenna shown here, must be bonded directly to the RFIC. This means there is no convenient connection. OTA measurement can be achieved with a horn antenna but lower signal levels and calibration are big issues.

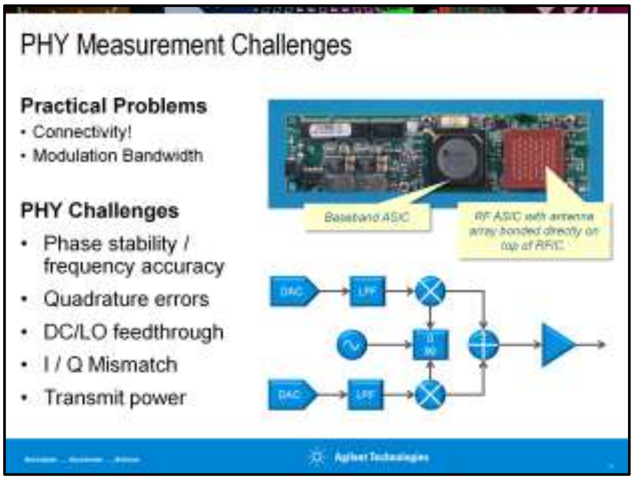

Before we get into actual measurements, it might be a good time to review some fundamentals of the topologies used for mmW circuits.

Here we have a generic I-Q modulation block diagram with DACs, Low Pass Filters, Local Oscillator and phase shifter, Modulator and finally amplifier.

As we mentioned before, at 60 GHz basic physical impairments present themselves in the form of

- phase instability
- Quadrature errors
- Skew
- mismatch
- and the like.

Background:
Test connectivity is a challenge at 60 GHz. The high operating frequency together with tightly coupled and integrated, multi-element antennas mean that **physical connection to the DUT will be difficult or impossible**. Even if a physical connection is possible, the electromagnetic performance of **the antenna** is dynamic, or at least varies with direction and polarization, and **cannot sensibly be omitted from a whole-device test**. So, for RF measurements, over-the-air (OTA) testing will likely become the norm.

Perhaps surprisingly, the big challenge is <u>not</u> 60 GHz. Obviously it's 10x higher frequency than current mainstream wireless technologies, but that's still familiar territory; the big challenge is modulation bandwidth. 2 GHz is <u>100x</u> wider than current Wireless LAN. New tools are needed,

What are the likely PHY problems at 60 GHz? Essentially the usual vector modulation challenges, just harder to solve. Imperfections that would be inconsequential at 5 GHz with 20 MHz of modulation bandwidth, can render a 60 GHz / 1.8 GHz wide signal incomprehensible.

1. **Phase stability / frequency accuracy** : Similar tolerances to other WLAN PHY's (20 ppm etc.) are required, but these are more difficult to achieve with large PLL multiplications from MHz references to GHz frequencies. Phase instability can also affect channel estimation.

2. **Quadrature errors** : An issue for analog IQ modulators, especially in direct conversion designs. As the carrier frequency increases, so too does the impact of small timing offsets between the “sin” and “cos” carriers on quadrature. Digital IF modulation is free of quadrature errors (being “mathematically perfect”) but requires much wider bandwidth D/A conversion and multistage frequency conversion to get to 60 GHz. Quadrature error can affect channel estimation.
3. **DC / LO feedthrough** : Unbalanced DC on the IQ inputs or LO leakage through the mixer can result in a non-zero carrier component in the transmitted signal, this will demodulate as an unwanted DC term (or a low frequency spur if carrier frequencies are not locked) which can impact EVM (if not removed) and can affect channel estimation.
4. **I/Q Mismatch** : Gain mismatches between I and Q distort the transmitted constellation compromising EVM, but the biggest IQ mismatch problem is timing, just a few 10’s of picoseconds of differential delay is significant at gigabit symbol rates. It is also important for low EVM that the baseband group delay is linear and the amplitude response is flat.
5. **Transmit power** : Wideband (2 GHz) linear power amplification is a challenge. Some vendors may support only BPSK / QPSK modes (which can use high-efficiency amplifiers), but QAM and OFDM require linear amplification. Amplifier back-off to avoid non-linearity reduces available power and hence range. Over-driving the amplifier to boost range will cause non-linear (3rd order) distortion and adjacent spectrum regrowth. The amplifier may also have turn on/off transients and amplitude stability issues.

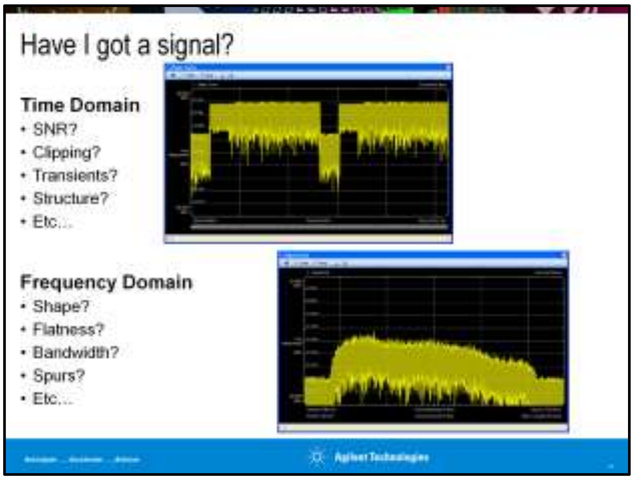

The first simple question is, “Have I got a signal at all?”.

We can use the time and frequency domains so make some first order assessment of the signal.

In the time domain, we can definitely see things like:

- Signal-to-noise ratio
- And we can identify some of the packet characteristics – for instance, we can see the preamble and then we can see some modulation peaks
- In the frequency domain, we are able to see some of the bandwidth as well as some of the shape or unflatness in the spectrum.

Background (not used in narrative)

For this purpose basic time (oscilloscope style) and frequency (spectrum analyzer style) displays can provide useful assurance. (The power versus time plot is also what you might expect to observe with a spectrum analyzer in zero-span mode. However, there aren’t any SA’s with RBW’s approaching the signal’s 2 GHz BW. These plots were made using a VSA (block converter and scope) with the ability to capture the full bandwidth of the signal and preserve the signal shape.

Even from these basic traces, with a little experience, you can often judge whether a signal “looks right”.

In the Time Domain

- What’s the signal to noise ratio (roughly)? - Because 11ad is a burst mode signal you can guesstimate the received SNR by “eyeballing” the mark-space ratio.
- Does the signal look clipped or compressed? - Especially in a QAM or OFDM signal we expect the header and payload fields to have crest factor >> 1.

- Are there any transient effects? - Amplifier issues such as slow power ramping or ringing may be visible.
- Can you see the transitions between the preamble, header and payload fields?
- Etc.

In the Frequency Domain

- Is the spectrum the right sort of shape? – If it's SCPHY, does it have the Gaussian / Cosine / custom shape you expect? If it's OFDMPHY is it the usual "sand castle" shape?
- Are there any big dips or peaks? – Some non-flatness is expected and will be equalized out, but extreme changes can defeat an equalizer and hence be problematic.
- Is the bandwidth about right? – i.e. about 1.8 GHz
- Are there any spurs or adjacent channel spectral regrowth? – . In band spurs will degrade EVM. "Clumps" of noise on either side of the expected signal may indicate intermodulation products due to non-linearities such as gain compression.
- Etc.

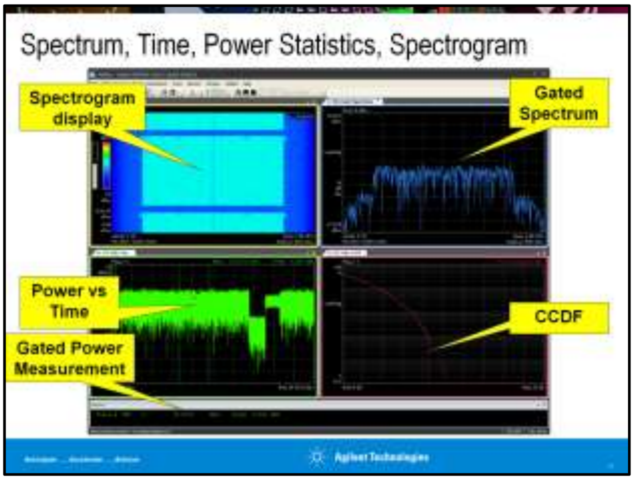

Here we can see the power of using vector signal analysis to analyze the signal – even prior to demodulation.

Here we see a spectrogram – Frequency in the X-axis, Time in the Y-axis and Amplitude in the Z-axis. And this can clearly point out that we have packets – which turn on and off in a burst – and we can also see a line up the center where the carrier is nulled.

Using time-gating markers in the lower left time display, we can measure power over a given interval – here a portion of the data field, the gated spectrum and its power statistics in the form of CCDF or cumulative complementary distribution.

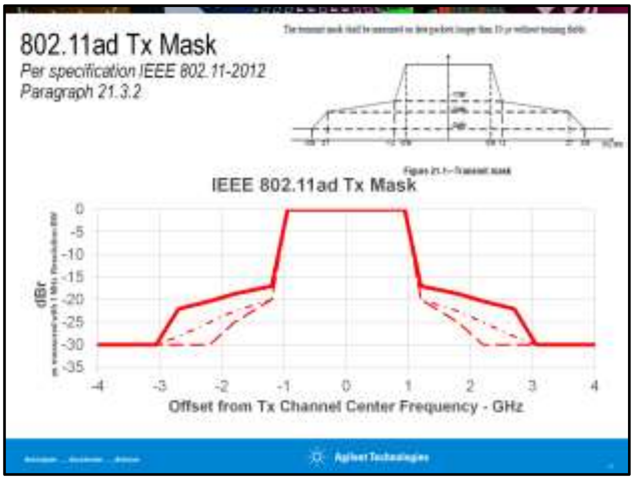

If we can look at the transmit spectrum.
The mask for 802.11ad started out as shown

<click>

but it was relaxed in two iterations to accommodate the spectral spread of signal carrier modulation and the reality of implementations

<click>

Background (not part of narrative)

If you're familiar with other OFDM specifications, you will recognize this shape.
But there are two aspects to note.
The first is that, because 11ad uses both single carrier and OFDM modulation, the breakpoints at -25dBr have been pushed out a bit to accommodate the rather more spread-out spectrum of single carrier modulation,
and second; even though a mask was approved by the ITU, it has been superseded by a subsequent draft and is now part of the released standard.
The mask was relaxed further so as to make the power amplifier designer's job a bit easier by permitting higher levels of out-of-band distortion products.
The spectrum mask is expressed in decibels relative to the signal level at the band center (dBr).
The maximum permitted transmitter power varies by country, but in general +10dBm can be taken as a practical limit.

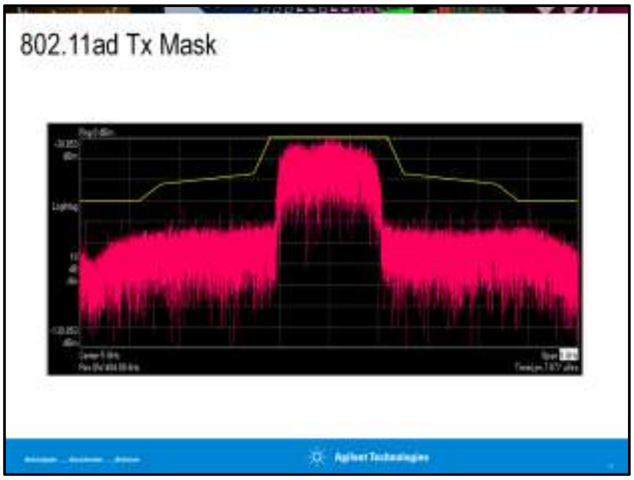

If we apply today’s mask to a real signal, we can see right away what margin and issues we might have.

Here we show the metrology-grade signal from Agilent instrumentation. There is lots of margin, but we can clearly see some spurs.

A product implementation would most likely be closer to the mask.

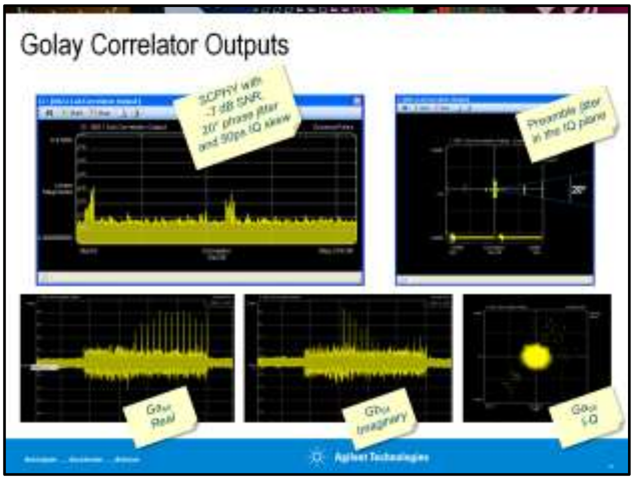

As you have probably realized from the PHY description, the Golay correlator outputs are a goldmine of information – a built-in diagnostic tool.

Before demodulating the 802.11ad header and payload, and especially the received signal that is so degraded that demodulation is impossible. We can learn much from looking at the correlator outputs.

When dealing with an unknown signal, it is often a useful first check to look at is the real component of Ga128 / Gb128 correlator outputs. On a very clean signal you will, of course, expect to see the pattern of preamble correlation spikes illustrated earlier, but it is remarkable that even when the signal is buried in noise and other distortions the preamble correlation will still be discernible, though you may need to switch to linear magnitude to remove phase rotation effects.

The trace shows the Ga128 output for an SCPHY signal suffering from a -7 dB SNR, 20° phase jitter and 50ps I/Q skew. If you cannot see any significant correlation spikes there is something fundamentally wrong with the signal, and probably the most common problem will be spectrum inversion due to frequency conversion.

Reaching 60 GHz performance at commercially viable prices is a challenge. In particular, it is difficult to achieve sufficient local oscillator phase stability. Phase noise will often be the single biggest impairment in a new design, so it is useful to know that the Golay correlators can help with early diagnosis of phase stability problems.

<click>

Ideally, correlator outputs will lie on the real axis, but impairments cause the spikes to rotate in the complex plane and acquire an imaginary component. This looks like amplitude modulation on the real signal, but can be more clearly seen as phase modulation if we plot the IQ plane.

This bottom trace is actually a real signal taken over the air, so we can see the clear rotation although it is not as clean in the I-Q plot as the manufactured signal we have up above, but we can see a discernable spread in phase noise.

Background: Not for narrative

The correlator outputs can be used to detect a vast array of behaviour, such as :

- unwanted frequency inversion,
- surface possible phase stability issues
- identify the basic PHY type
- confirm the presence and timing of SCPHY guard intervals and BRP fields
- even to perform a rudimentary decode of the CPHY.

These can all be useful diagnostics if an apparently “clean” signal is proving difficult to demodulate.

Looking at the magnitude of correlator output s is also useful, especially for detecting variations in PA gain at the start of the burst, or for detecting missing peaks from improper PA gating.

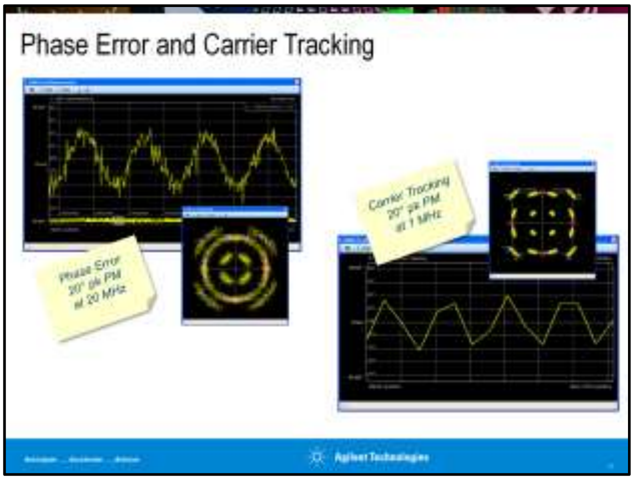

We can use the Error Vector phase error information to see any patterns in the phase domain – here a higher rate phase error.

If we engage Carrier Tracking, we can determine if there are any carrier frequency deviations during the packet duration and compensate for them. This would be true for a slower moving error mechanism.

The resulting correction trace will give us some indication of the nature of the deviation.

Here we see phase modulation at a 1 MHz rate.

Background: (Not part of narrative)

Practical receivers will include carrier and phase tracking loops to remove those instabilities from the received signal, however it is still good practice to minimise these impairments in the DUT where possible, to reclaim that performance margin for the link budget.

In a measuring receiver it is useful to display the carrier and phase tracking error versus time to help quantify the nature of the perturbation; in particular whether there is any coherent content in the error signal.

In the graphics we see on the left, a phase error plot revealing that the signal is subject to phase modulation at 20 MHz with a peak amplitude of 20° . On the right, we see a carrier tracking plot (for a different signal) revealing that it is subject to phase modulation at 1 MHz with a peak amplitude of 20°. **E**ach data point represents the location of a GI symbol.

The best plot to look at depends on the nature of the interference, a 1 MHz modulation is just within the bandwidth of the carrier tracking loop so we see variation in the carrier tracking plot, but at this frequency the tracking is not perfect so there is some residual phase modulation on the constellation. At lower perturbation frequencies the displayed constellation would be free of phase modulation, but we would need a longer time record to compute the frequency from the carrier tracking information.

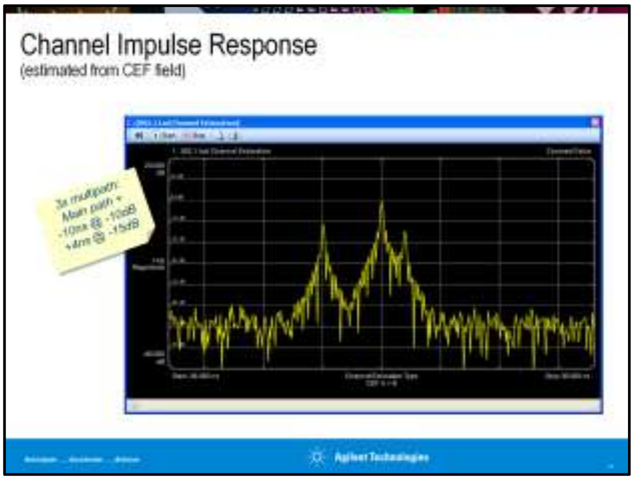

Channel estimation gives a very clear indication of the multipath content of the signal. In this example we have contrived to have three clearly distinct paths for illustrative purposes, but it is more common for there to be a “smear” of delayed responses in the 0 to 20ns range.

You will note that the center X-axis is <u>zero</u> time, and therefore the peak to the right has come later and is attenuated about 15 dB. But we also notice a peak to the left which a little bit earlier and is suppressed about 10 dB. But it is occurring in negative time. Now how can that be?

This correlator output will essentially take the maximum response and place it at the center of time in the display. So any responses to the left will have actually arrived earlier, but be weaker. So this is a case where the peak response was actually NOT the first one, or the direct one. In fact none of them could be direct, the very first response was attenuated about 10 dB. The best response was a little bit later and 10 dB higher. So we could make an assumption or a guess that the one to the left is actually the main, direct transmission and that the one with the highest peak was less obstructed, but arrived a little later. And then the one to the right got there even later and was attenuated as well.

Background (not part of the narrative):

The data for this measurement comes from processing the channel estimation field (CEF) in the preamble. The CEF contains two estimation opportunities which may be examined separately (in our nomenclature CEF A and CEF B) or combined as CEF A+B as illustrated here. Which display option is more useful will depend on how quickly the channel is changing.

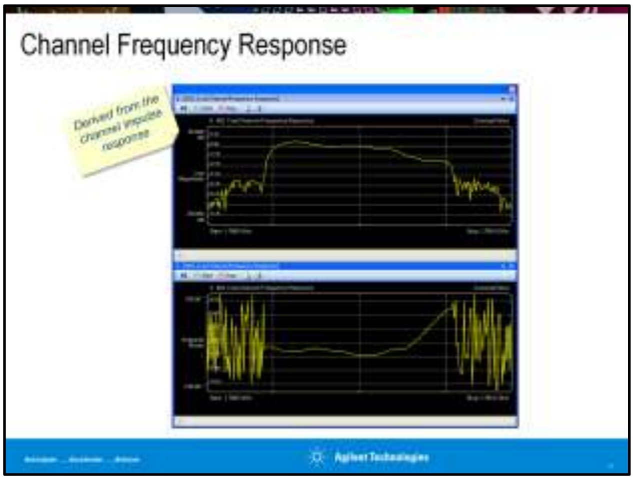

Now, if we look in the frequency domain, we can get the channel equalization. Like the DUT, the measuring receiver includes a channel equalizer to compensate for the non-ideal amplitude and phase response of the received signal.

In performing this correction the equalizer necessarily estimates the “error” to be compensated, which is actually an estimate of the channel frequency response, and that is the basis of this display.

It is an estimate, because the equalizer may not be able to exactly compensate for the effects of the channel if there are extreme changes in the response, but in general it should match the in-band part of the corresponding spectrum display well.

Equalizer design in a measuring receiver is always a balancing act.

On one hand we want a robust implementation to ensure measurement integrity, but on the other hand we do not want to implement an equalizer with disproportionately better performance than the equalizer that is likely to be in the DUT because that could mislead a test engineer into thinking their marginal signal was actually acceptable.

We should now mention that the channel response includes the transmit shaping filter response because the Golay codes pass through both the transmit filter and the channel.

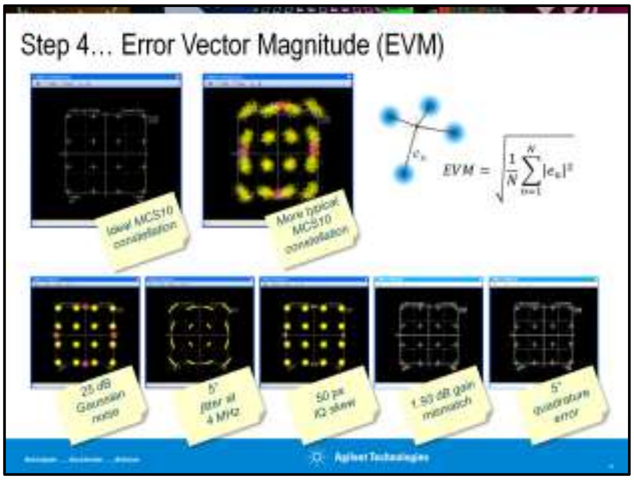

If there is a single, catch all, figure of merit for a complex modulated signal it has to be Error Vector Magnitude, or EVM.

I expect most of you are familiar with the concept of EVM. A complex modulation is constantly varying in amplitude and phase versus time. Its complex value is sampled periodically at the symbol rate and, at those sampling instants, a perfect signal would be found to be exactly at one of a finite number of pre-defined positions in the complex plane, which can then be uniquely mapped back to an equivalent data pattern.

In practice a real signal will rarely be at exactly the ideal value. We hope it will be sufficiently close to the ideal that it can be unambiguously decoded, but its actual position will differ from the ideal by an error vector. If we compute the average magnitude of these error vectors, we obtain an objective measure of the signal perturbation.

In a real system, this perturbation is an aggregation of many effects; some, such as noise, phase jitter, inter-symbol interference, timing skew, and symbol clock frequency error all smear or diffuse the received sample values relative to the ideal in very distinctive ways, cumulatively creating the familiar "fuzzy ball" round each ideal position.

Other impairments such as IQ gain imbalance, quadrature error, IQ DC offset, and gain compression distort the geometry of the received signal so that even a best fit to the ideal constellation necessarily leaves some points "off target", again degrading EVM.

There is an art to examining an I-Q constellation and determining what impairments exist, but it can be done and is quite useful.

Objective summation of the error vectors to give a single EVM result is a useful benchmark, but the subjective assessment of an impaired constellation based on an understanding of these error mechanisms can often provide greater insight to fault-finding the root cause problems.

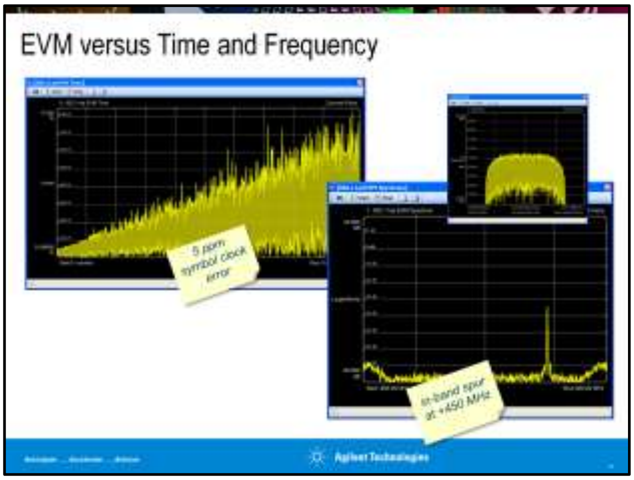

So EVM when considered as a single "figure-of-merit" result is a relatively blunt instrument.

However, each single EVM value can legitimately be considered as a sample of the continuous error signal that represents the difference between the ideal and actual signals. As such, this time-sequence of EVM values contains information about the time and frequency domain properties of the error signal.

EVM Time obviously shows how EVM varies over time and can often quickly identify drift or instability in a signal. It can also identify amplitude dependent effects such as if the EVM degrades when switching from a QPSK header to a QAM16 payload, suggesting that perhaps there is a gain non-linearity.

This EVM vs. Time example shows a distinct accumulation of error over time. This is an indication of symbol clock error because we are accumulating error as we get further and further away from the actual symbol clock.

EVM Spectrum is simply the Fourier Transform of that time signal. If there are any dominant frequency components in the error signal they will show as increased energy at those frequencies in the spectrum plot. In this case we can see a spur, which we could not see under an in-channel spectrum measurement.

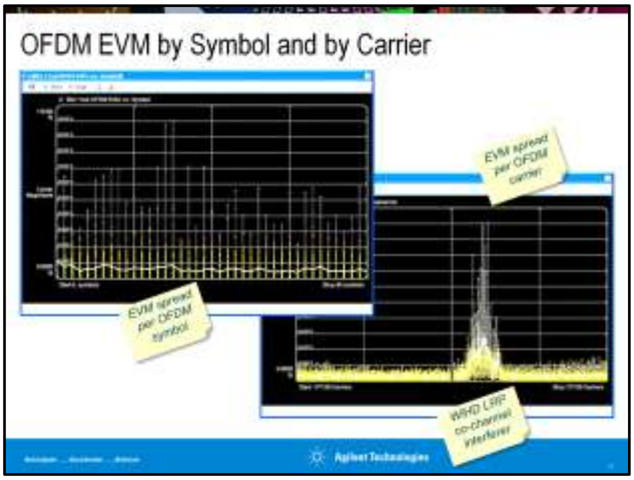

A variation of the EVM time and EVM spectrum display is possible for the OFDM modes.

In this display we can again see the variation of EVM over time and frequency, but this time tailored to the OFDM modulation structure by showing the EVM for each symbol and for each modulated carrier of the OFDM signal. The colour coding shown, matches that of the equivalent constellation display. You can see the position of the pilot tones as the periodic white lines in the spectrum.

For each x-ordinate, the y-ordinate data shows the spread of EVM values for that (symbol / tone), and the white trace is the mean value.

In this case we see that there is a modulated in-band interferer. In this case, it is a WirelessHD LRP burst at Fc + 159.625 MHz at about 25 dB down.

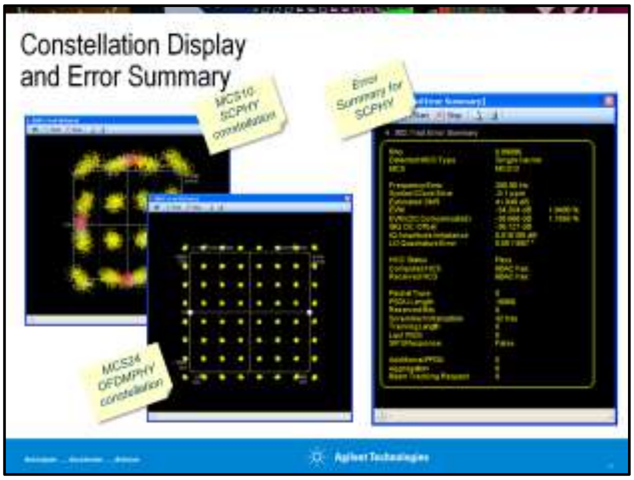

We already discussed the I-Q constellation.

A Modulation Quality Metrics display gives us EVM and other parameters we can extract from the measurement, such as frequency error, DC Offset, quadrature error, IQ Gain imbalance, etc. As well as reading the header information.

Background: (Not part of narrative)
The 81199A provides a tabulation of the key modulation quality metrics.

Rho : A figure-of-merit on the preamble, somewhat analogous to "sync correlation" in LTE. It is computed from the correlator outputs and is intended to indicate the fidelity of the preamble. A Rho of 1.0000 indicates a "perfect" preamble.

Frequency Error : The estimated frequency error in the received carrier frequency (subject to the accuracy of the measuring receiver's own frequency reference). This is a signal-based estimate so there is usually a few hundred Hz of residual "noise" on this measurement, but significant offsets (10's of kHz) are accurately reported. The measurement range is approximately $F_{c(nominal)}$ ± 6.5 MHz.

Estimated SNR : The estimated random content (variance) of the received signal (additive noise and phase noise, for example). It's main value is as a comparison with EVM, if the two numbers are similar magnitude (ignoring the sign difference due to how the measurements are defined) then you can assume that random noise is the dominant component of the EVM result. If there is a significant difference then there are likely other impairments present. At very low SNRs this estimate can become inaccurate.

EVM : The primary modulation figure of merit, as already discussed.

EVM (DC Compensated) : IQ DC offsets will often be removed by the DUT receiver, so it is useful to also compute EVM without this contribution, to give a more accurate assessment of the signal quality as it will be seen by the DUT.

IQ DC offset : The DC content of the signal.

IQ Amplitude imbalance and LO Quadrature error : The estimated value of these impairments, however we should take a moment to discuss a potential difficulty in correctly evaluating these parameters. This is picked up on the next slide.

HCS Status: Finally we take advantage of the Header Check Sum (HCS) in the 802.11ad header to confirm that at least the header data has been successfully (or unsuccessfully) demodulated.

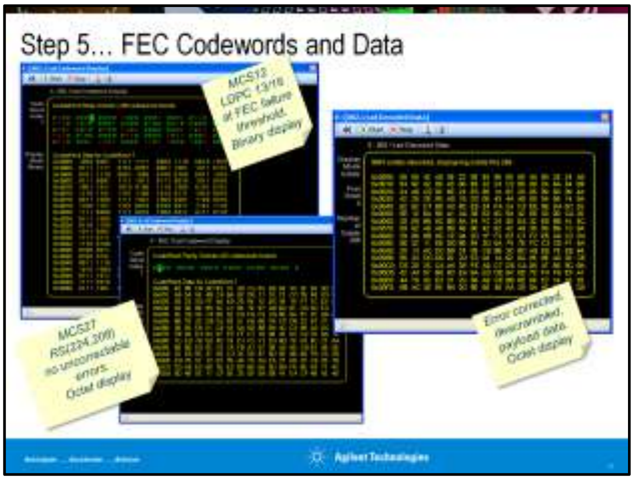

As long as we go that far, we might as well decode the data.

Here are codeword displays to tell us if error correction could recover from the errors and the decoded bits themselves.

Background: (Not part of narrative)

In order to extract the parametric results discussed so far, a measuring receiver has to implement a more or less complete demodulator, including the error correction algorithms to recover, as far as possible, the original transmitted data for reference purposes. These decoding processes can also yield useful diagnostic data.

As described earlier, the data is grouped into LDPC or Reed-Solomon codeword blocks for error correction. All block error correcting codes have threshold beyond which they can no longer correct all the errors present. At this point the error correction fails; LDPC iterations fail to converge, RS error locator and magnitude polynomial calculations fail to solve. Given the strength of the codes used, this threshold is typically at such a high error rate that when it fails, the failure is catastrophic; the recovered data goes from clean to garbage in the space of a very small SNR change; the so called “cliff effect”.

So it’s useful to know when error correction failures are starting to occur.

It can also be useful to know where in the packet they start to occur first.   Early in the packet? Perhaps there is a turn-on transient issue, what can we see in the correlator traces? Late in the packet? Perhaps there’s a drift issue, what is EVM time showing?

Or perhaps you just need to verify that your LDPC encoder design is operating correctly in all its various ratios.

Finally we have the recovered payload data. Any known pattern (including all 1's and all 0's) can serve as the basis for a packet BER calculation. 8 or 32-bit count data is easier to "eyeball", but PN data is usually considered the most representative of real data traffic.

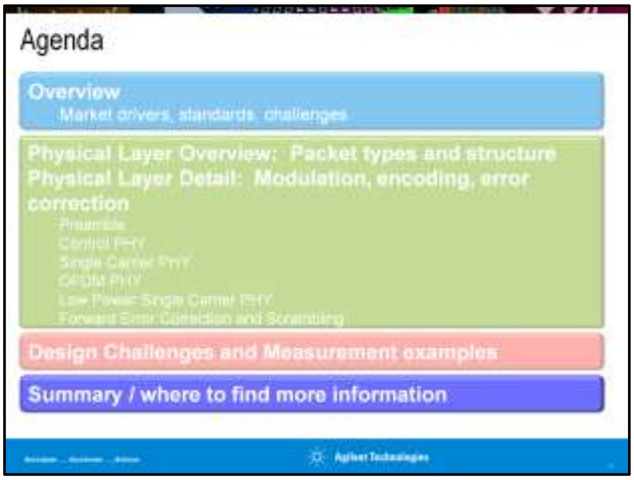
Agenda
Overview
Market drivers, standards, challenges
Physical Layer Overview: Packet types and structure
Physical Layer Detail: Modulation, encoding, error correction
Preamble
Control PHY
Single Carrier PHY
OFDM PHY
Low Power Single Carrier PHY
Forward Error Correction and Scrambling
Design Challenges and Measurement examples
Summary / where to find more information
Agilent Technologies

In summary…

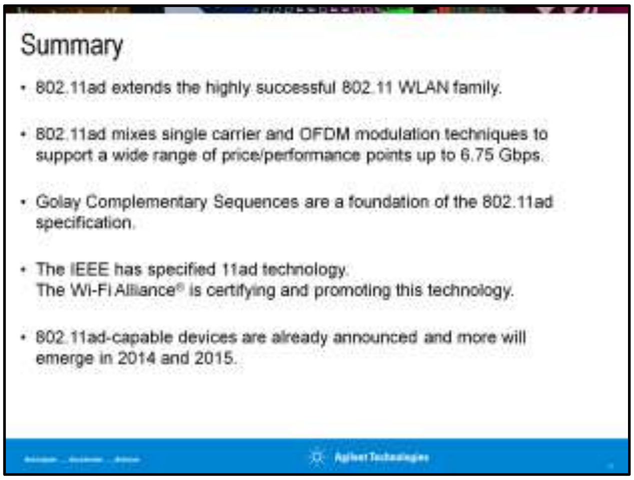

We've seen that 802.11ad is an extension of the highly successful 802.11 Wireless LAN family

composed of single carrier and OFDM modulation types

employing Golay Complementary Sequences in a variety of roles.

The Wi-Fi Alliance will certify the 60 GHz MAC and PHY with associated certifications for adjoining serial interfaces to be determined

And we should see more 11ad-capable products in the coming years.

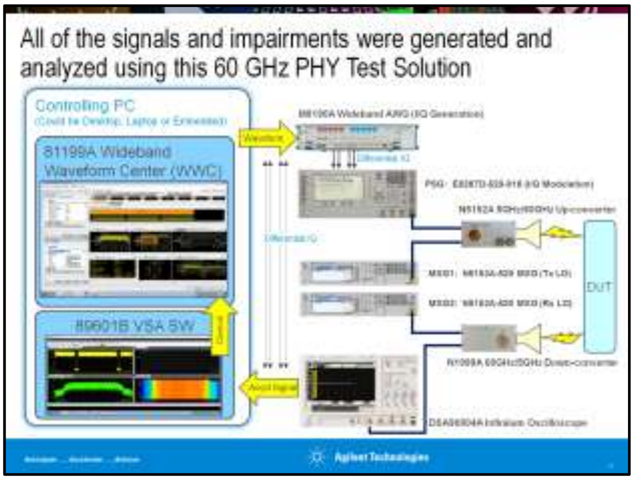

All of the measurements shown in today's presentation used this setup and solution from Agilent's 802.11ad PHY Test System.

We would be happy to go into details of the system during a break or as a follow-up.

Background: (Not part of narrative)
Waveform data is downloaded to an M8190A wideband Arbitrary Waveform Generator (AWG) where they are "played" to produce analog differential or single-ended I and Q signals representing the 802.11ad signal as a baseband (zero IF) modulation.

The I and Q signals can be fed directly to a DUT or, as in this configuration, fed to the wideband IQ inputs of an E8267D PSG Vector Signal Generator which can modulate the baseband signals onto an IF carrier in the range 4 to 44 GHz. This IF can be fed to a DUT if appropriate.

Shown here an IF frequency of 5 GHz is chosen, then the signal generator output can be fed to the IF input of an N5152A 5 GHz/60 GHz Broadband Up-Converter to produce a 60 GHz RF stimulus. An N5183A MXG Microwave Signal Generator provides the L.O. for the up-converter to determine which of the 4 RF channels the final modulation is delivered to. The L.O. frequency is simply determined by the equation $F_{LO} = (F_{CHAN} – 17.5\ GHz) / 4.0$.

The output of the N5152A Up-Converter is a WR-15 waveguide bulkhead to which a standard gain horn (SGH) can be attached for over-the-air stimulus of the DUT. Alternatively a V281A WR-15 to 1.85mm adapter can be used to implement a cabled connection to the DUT. The same connectivity options are available on the receive side because the N1999A 60 GHz/5 GHz Down-Converter also has a WR-15 waveguide bulkhead as its input.

The Down-Converter is tuned to the RF channel of interest by a second N5183A MXG Microwave Signal Generator according to the equation $F_{LO} = (F_{CHAN} – 5\ GHz) / 4.0$. The 5 GHz modulated IF output from the Down-Converter can be digitized directly by the DSO/DSOX 90000 series Infiniium real-time oscilloscope. Depending on the chosen model of Infiniium 'scope, IF's up to 30 GHz can be acquired.

Alternatively, the oscilloscope can be configured to accept differential or single-ended IQ signals directly from a baseband DUT if appropriate.

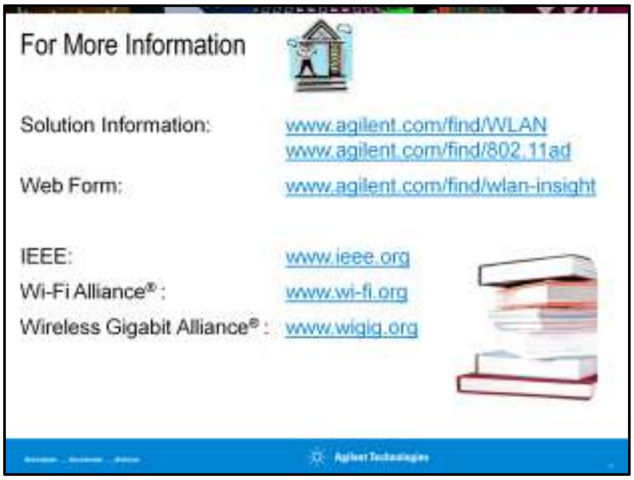

We have only scratched the surface of this WLAN specification.

You may find additional information on the Agilent website at w w w dot agilent dot com slash find slash 802.11ad.

There you will find a link to an in-depth tutorial.
Here are the URLs for the IEEE, WiGig Alliance and Wi-Fi Alliance, respectively, for your convenience.

Just a reminder that on November 1st, Agilent Test & Measurement will become Keysight Technologies.

With that, are there any additional questions?