Validating and Troubleshooting OFDM

Considerations and Approaches

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Ben Maxson
Product Engineering
Copper Mountain Technologies
3905 Vincennes Road, Indianapolis Indiana 46268
+1.317.222.5400
ben.m@coppermountaintech.com

Section 1: Introduction to OFDM

Subscriber demand for higher data rates has grown relentlessly since the introduction of DOCSIS® more than 15 years ago. Service providers using DOCSIS® to deliver data services over hybrid fiber-coaxial (HFC) networks must always be looking for methods of delivering higher data rates to customers, while at the same time maintaining reliability and controlling maintenance costs.

Historically improvements in data service robustness and speed have been obtained through incremental enhancements to DOCSIS®-compliant customer premises equipment (CPE) and cable modem termination system (CMTS) equipment, through ongoing optimization of the many configuration parameters of the DOCSIS® network, through node splitting, through systematic ingress control efforts, and most recently through bonding of 4, 8, 16, or even more quadrature amplitude modulation (QAM) channels, in both the upstream and downstream.

All these advancements have been built on the same physical layer modulation techniques: single-carrier QAM (SC-QAM) with Reed-Solomon (RS) Forward Error Correction (FEC). Meanwhile, advancements in semiconductor feature size, cost, and performance have enabled dramatic performance improvements in other sectors of communications technology, most notably multi-carrier modulation including Orthogonal Frequency Division Multiplexing (OFDM) and advanced FEC including Low Density Parity Check (LPDC) codes.

OFDM is a modulation technique which divides the available channel bandwidth into many smaller narrowband channels called subcarriers. OFDM has several advantages over traditional single-carrier or bonded SC-QAM modulation. In this paper, following a brief introduction of OFDM technology as compared with SC-QAM modulation in the HFC forward path, we focus on the implications of OFDM modulation for test and measurement applications. OFDM has also been chosen for inclusion in the forthcoming DOCSIS® 3.1 specification. For simplicity, in this paper we focus on published standards for downstream OFDM, though some of the same principles will apply in part or in full to the upstream and to DOCSIS® 3.1.

Mass deployment of consumer technologies utilizing OFDM, including WiFi, MoCA, DVB-T/T2 and many others is well underway. OFDM modulation has been applied to cable networks, first via waveforms designed for terrestrial broadcast such as DVB-T/T2, DTMB, and ISDB-T—which some cable operators are obligated to carry on their HFC networks—and more recently as DVB-C2 which has been specifically designed for the HFC medium.

Specifically, we first consider the implications of OFDM for some of the most popular digital performance test metrics, such as modulation-error ratio (MER), bit-error ratio (BER), constellation, and equalization, then offer some approaches for test and measurement of OFDM based on past experience and future possibilities.

Section 2: Motivations for OFDM

OFDM divides the available bandwidth of a channel into multiple narrowband digital subcarriers with sufficient frequency separation to prevent them interfering with one another [1]. Each carrier can employ a different QAM modulation format and different error correction parameters. Bandwidth efficiency can be maximized by choosing these parameters according to the predicted signal to noise ratio (SNR) or measured SNR(s) at the receiver(s) of the data being sent [2].

An OFDM-modulated channel is somewhat analogous to the cable channel plan as it exists today. Nominally, the approximately 800 or 900 MHz of available downstream spectrum is divided into non-overlapping channels. Each channel contains has a SC-QAM carrier with a certain bandwidth, nominally 6 MHz or 8 MHz. The channels avoid interfering with each other because they are separated in frequency, by their occupied bandwidths plus a small guard band, specifically.

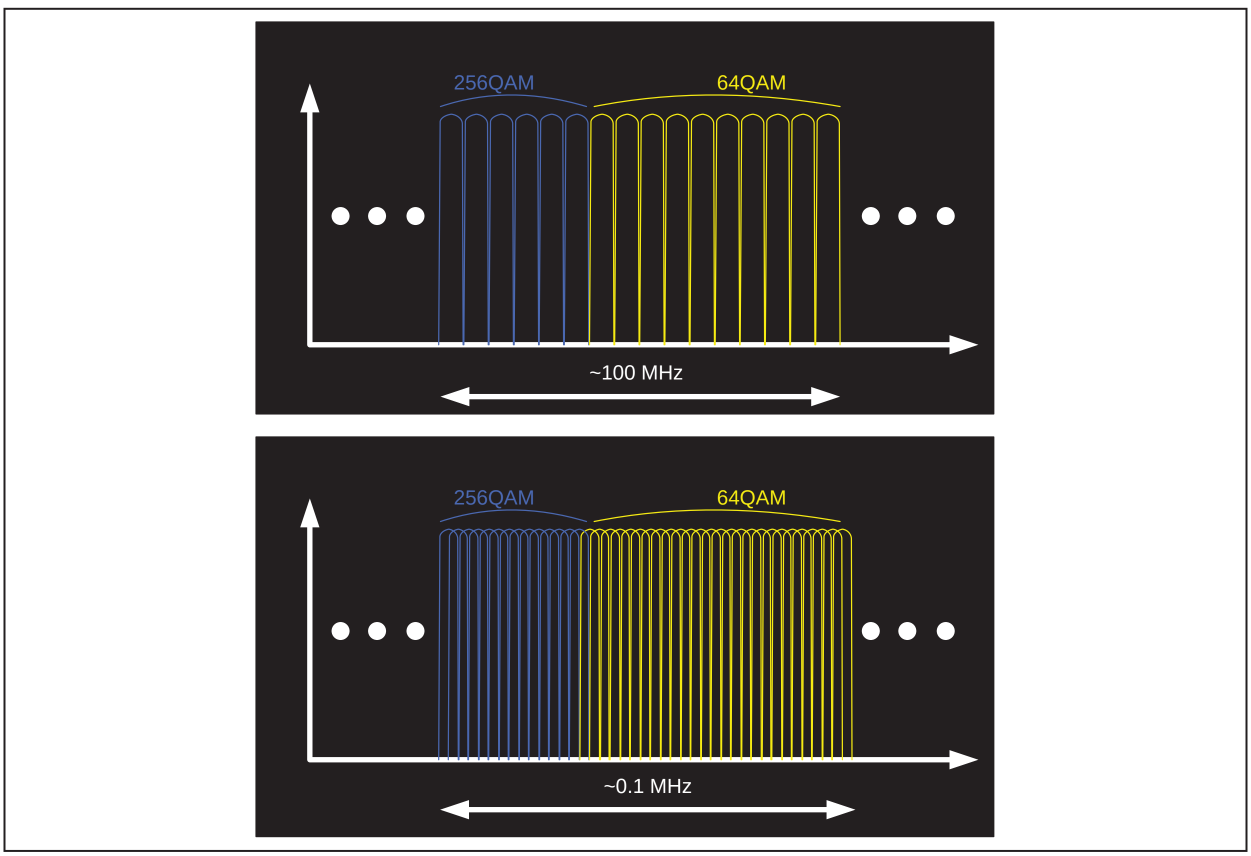

Figure 1 – An OFDM channel is similar to the downstream of a cable network; multiple carriers with a diversity of QAM modulations are packed tightly together, but separated sufficiently to avoid interfering with one another. In a cable downstream, 120 or more carriers are present with bandwidth 6-8 MHz. In an OFDM channel, thousands of carriers occupy a single 6-8 MHz channel.

OFDM is conceptually similar to the downstream channel plan except that whereas each SC-QAM carrier occupies 6 MHz or 8 MHz, each OFDM subcarrier occupies only a few kilohertz, 2 kilohertz for example. It can be mathematically shown that if the transmitter satisfies rules relating to pulse shaping and time synchronization, the carriers can be overlapped in frequency without causing interference between them. As a result, carriers can be packed very closely together, doing away with the overhead of guard bands in SC-QAM networks. In a spectrum analyzer, adjacent modulated carriers of OFDM cannot be resolved, they appear as a noise-like haystack.

Feature	SC-QAM	OFDM [2]
Carrier spacing	6 – 8 MHz	~2 kHz
Carriers per 6-8 MHz	1	Thousands
Carriers in downstream	Up to 150	Up to hundreds of thousands
Symbol rate	~90% carrier spacing	~100% carrier spacing
Inter-carrier guard band	~400 kHz	None

Table 1 – Similarities and differences between a SC-QAM modulation and OFDM modulation.

Why go to the trouble of using OFDM, with the complications of separating data into multiple carriers only to recombine the data at the receiver? One primary reason is ease of equalization [3]. Since each carrier is narrow in frequency and long in time, it is very easy to equalize. Wireless channels, like HFC channels, can have significant variation in gain or group delay over a wide span such as 6 MHz or 8 MHz. But if one “zooms in” on even a fairly bad channel to a span 2 kHz, the channel will appear to be flat in amplitude and group delay. An equalizer to compensate such a channel is trivial; it can be accomplished with a single, complex-valued tap to accomplish gain correction and phase rotation.

Another benefit of using numerous narrowband carriers packed together is this allow for customization of the modulation on each carrier according to its predicted or measured SNR [3]. Again, an analogy can be made to the full cable plant, where one might operate 256QAM channels at lower frequencies or slightly higher levels, i.e. when one can be more confident signals will arrive at CPE with sufficient SNR for successful demodulation. Or one might choose to operate more robust 64QAM channels in the FM band, where they are more likely to encounter interference from ingress at some endpoints.

OFDM extends this flexibility to customizations within an individual 6 MHz or 8 MHz channel, so one need not decide on a single modulation format for the entire channel. If for example the upper portion of a particular channel is known or found to overlap with LTE ingress and therefore be unsuitable for 256QAM, we can run 64QAM on that portion of the channel, and 256QAM on the remainder. In DVB-C2, a variety of modulation formats and error correction rates are available from 16QAM to 4096QAM and 2/3 to 9/10, respectively, suitable for SNR between 11 dB and 35 dB [2].

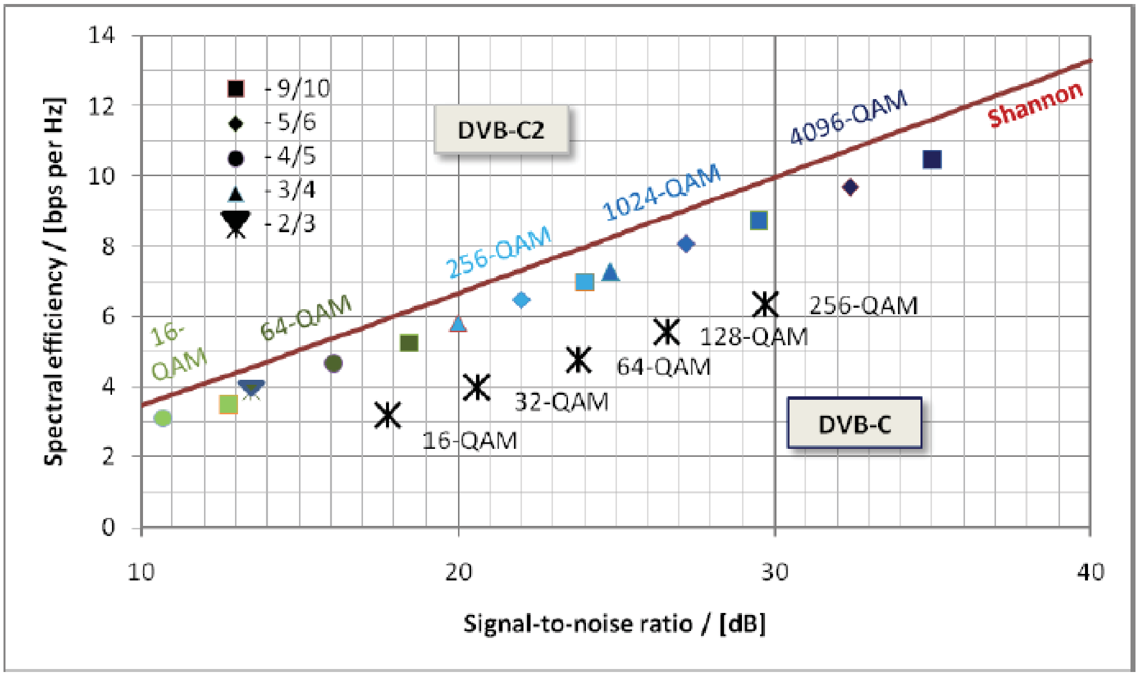

Figure 2 – DVB-C2 specifies 13 different QAM/FEC configurations which allow efficient utilization of channels with between 11 dB and 35 dB SNR [2].

In fact, one can dynamically decide which modulation format to use for each subcarrier based on the actual received quality as reported by each receiver using the channel at a given time. Such feedback, known as Channel State Information, eliminates guesswork with regard to which modulation parameters are ideal for a given receiver or group of receiver. Instead, these can be optimized automatically and dynamically [3]. This technique is called adaptive modulation and it has the potential to match the quality of services provided to the quality of the channel available, providing for extremely robust service delivery.

Section 3: OFDM MER

In the previous section we introduced OFDM modulation and described some of its main benefits relative to conventional, single-carrier QAM modulation. Let's dive in a layer deeper to the demodulator algorithms inside an OFDM receiver, since after all, the test results ultimately displayed to cable test and measurement users are usually derived directly from demodulator chipsets.

The OFDM receiver's processing is similar to the SC-QAM receiver processing in many respects. At a high level, the demodulator must first estimate and remove frequency offset between the transmitting modulator and receiver's tuner. Likewise, a symbol timing clock offset must be determined and compensated. Finally, the phase and amplitude variations of the channel must be removed through equalization.

The received signal is now ready for symbol detection, or slicing, where the soft symbol is converted to a received symbol according to which ideal symbol is most nearby in the constellation diagram. The amount of error between the ideal received symbol and actual received symbol, the Error Vector, is sampled and averaged to compute an RMS Error Vector Magnitude, from which the Modulation Error Ratio can be readily found. MER is used as a proxy for the SNR of the SC-QAM channel; to the extent an ideal or nearly ideal transmitter and receiver are used, this approximation is valid.

In a SC-QAM measurement device, there is one MER for each channel in the downstream, so that no more than 150 values are needed to characterize 900+ MHz of digital carriers' health. But in a network using OFDM, each narrowband subcarrier has its own MER [4]. Hundreds of thousands of MERs are required to represent a downstream full of OFDM signals. Determining an appropriate presentation format for the MER measurements at a single network endpoint or test point is obviously a challenge.

One possibility is to plot MER versus subcarrier index in graphical format [4]. Variations in MER may be due to variation in receive level or ingress strength among the subcarriers, or a combination of both. But with the number of MERs present in an OFDM downstream, it is impractical for measurement users to review the MER of every sub-carrier. This type of measurement display is more relevant to troubleshooting of a single 6 MHz or 8 MHz channel known to have performance issues, much in the same way measurements showing the noise floor under SC-QAM channels are used today.

Another possibility, and a common approach taken in OFDM test receivers, is to simply compute the average MER across all carriers in the channel [5]. This approach will result in a similar result to that of a SC-QAM channel MER operated in the same network at the same frequency. Drawbacks of simply averaging the subcarriers' MERs include the fact that a channel with noise highly concentrated in a small frequency region—with high SNR elsewhere—may have the same MER as a channel with moderate SNR at all frequencies, but this distinction may be useful to inform the troubleshooting process followed by a measurement user.

Another possible approach is to provide statistical analysis of the MER error vectors to the user, such as minimum MER, maximum MER, and average MER taken across all carriers. Drawbacks of this technique include risking confusing the user with more data than is actionable. Consider one network endpoint affected by narrowband ingress at a strong level; its minimum MER may be quite poor, but service delivery completely unaffected because the FEC and adaptive modulation in use can fully overcome the impairment. In fact, in general when adaptive modulation is used and is working correctly, the modulation format employed by the transmitter will have been selected to overcome the noise and other impairments of the channel, so that poor MER may not correspond to non-zero BER, let alone a customer-impacting degradation.

With conventional QAM modulation, it is not too difficult for a technician with some training to remember what a good MER is for 256QAM and 64QAM (35 dB and 30 dB

Figure 3 – OFDM test receivers can display MER versus subcarrier in graphical format. Here the MER of a 64QAM DVB-T channel is shown for subcarriers 0-6816. MER ranges from 37 dB to less than 25 dB, perhaps due to a strong micro-reflection in the network.

perhaps, respectively). However with 13 modulation/FEC formats, one now will have 13 such numbers to memorize instead of two, not included the common differences in requirements among MERs expected in the headend, network, and at the CPE. Thirty-nine parameters is too many to be memorized and used in everyday work.

QAM modulation	FEC rate	SNR @ 1E-6 + 6 dB margin, dB
4096	9/10	40
4096	5/6	38
1024	9/10	36
1024	5/6	33
1024	3/4	31
256	9/10	30
256	5/6	28
256	3/4	26
64	9/10	24
64	4/5	22
64	2/3	20
16	9/10	19
16	4/5	17

Table 2 – DVB-C2 QAM modulations, FEC rates, and SNRs required for 6dB margin to virtually error-free operation. 30 dB SNR is required for robust 256QAM at 9/10 FEC. In comparison, SC-QAM requires approximately 35 dB MER for similar performance [2].

This all suggests that—for reasons of simplicity and prioritization both—we may wish to use another metric besides minimum MER or average MER, something that more closely approximates the current or potential service quality impacts of a particular reception scenario.

One such metric relevant to adaptive modulations is Bandwidth Efficiency as a percentage of theoretical or target efficiency. This approach has been taken with respect to MoCA communications for example, which do use adaptive modulation to overcome noise and signal level issues. The total available bandwidth of a particular MoCA link can be computed by finding the modulation profile being employed by each subcarrier and comparing the total throughput available to the theoretical maximum. That is, the ratio of the link's actual maximum data rate to the data rate which would be available if all carriers utilized the highest modulation order and least FEC available within the MoCA standard [6].

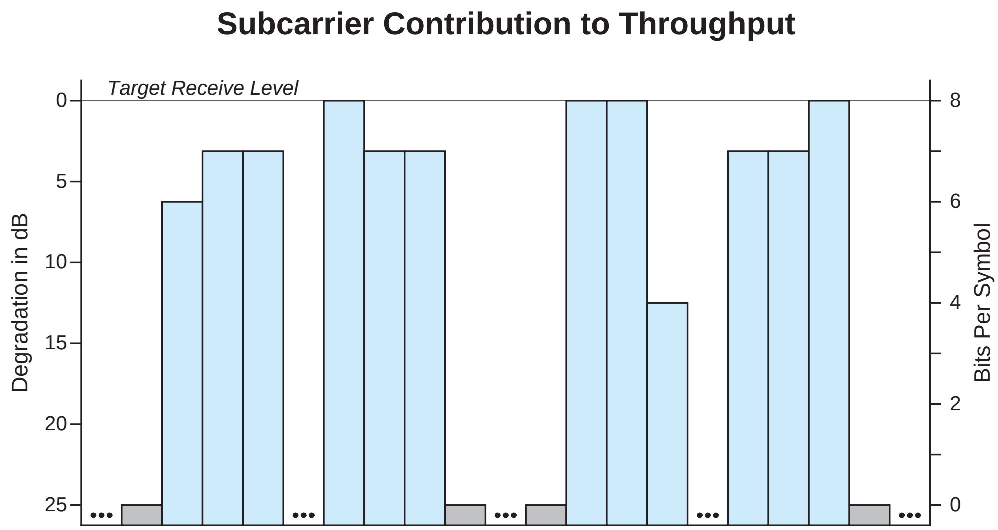

Figure 4 – In some test devices, the QAM modulation supported by the channel—and to which the MoCA link will adapt—is predicted for each subcarrier based on its receive level relative to a target level [6].

Throughput as % of Maximum	**Quality Score**
<= 50%	**0**
60%	**2**
70%	**4**
80%	**6**
90%	**8**
100%	**10**

Table 3 – The achieved throughput expressed as a percentage of the maximum theoretical value can be mapped to a unit-less Quality Score. This is one method of providing a simple, intuitive performance metric for communications systems employing adaptive modulation [6].

Section 4: OFDM BER

Since OFDM is essentially comprised of many narrowband subcarriers, it is more configurable than a conventional single-carrier waveform. Configurability can be considered as belonging to two separate categories, the first of which—employing a mix of modulation types across carriers in a single channel—has already been discussed in the preceding sections. The second type of configurability pertains to mapping—via multiplexing and de-multiplexing—the data sent on the various subcarriers into one or more data streams to pass up the network stack. In ISDB-T and DVB-T2/C2, all data of a particular Layer or physical layer pipe (PLP), respectively, must be carried with the same modulation format (QAM type and FEC parameters). There are nevertheless many ways to tailor the mapping between subcarriers and one or more data streams [2].

In ISDB-T a group of subcarriers is called a Layer and a maximum of three layers (minimum of one) can be used per 6 MHz channel [7]. The number of subcarriers assigned to each layer is adjustable in increments of 1/13^ths^ of the channel called Segments. As few as 1/13^th^ the carriers and as many as 13/13^ths^ can be assigned to each layer. In ISDB-T the first Layer (Layer A) has the special, optional capability of being confined to the center 1/13^th^ of the channel, e.g. not frequency interleaved with the higher layers [7]. This capability is widely utilized in broadcast applications to provide a robust low bitrate services via low order modulation (QPSK normally) to low-cost, low power, narrowband receivers in cellular telephones. This capability is often referenced with the shorthand name “1 Seg”.

DVB-C2 generalizes the mapping between frequencies and data streams. Up to 255 separate PLPs—each with its own modulation+FEC format—can be combined in a single data slice (channel, nominally) [2]. Conversely, up to 255 data slices can be employed to deliver a single data stream, a capability akin to channel bonding in DOCSIS® 3.0. Frequency interleaving is only applied across PLPs within a data slice, never across data slices. A single stream can thereby easily be split across data slices, each of which is confined to a specific frequency range, then recombined at the receiver via PLP bundling, again conceptually similar to channel bonding.

In the context of OFDM test and measurement, this flexibility of allocation has several implications. First, within a given 6 or 8 MHz channel, there may be multiple (up to 3 in ISDB-T, up to 255 in DVB-C2) Layers/PLPs present, each of which may have a unique modulation format, FEC rate, and BER. As a result, whereas a given 6 MHz or 8 MHz SC-QAM channel is characterized by a single pair of BER performance metrics (Pre- and Post-FEC) DVB-C2 based OFDM networks potentially have many more. Second, depending on the frequencies allocated to each Layer or PLP of the channel, the differences in BER may correspond to differences in the degree of ingress at those frequencies or signal level variations. Just as with DOCSIS® 3.0, service quality problems may be attributable to any of several channels, potentially channels widely separated in frequency.

Finally, it is important to know many DVB-C2/T2 and ISDB-T demodulators are designed to demodulate only one logical channel (typically meaning one Layer/PLP) at time, to reduce complexity and power consumption. In normal applications of these chipsets, such as set top boxes and USB tuners, the user is only interested in viewing or recording one stream per demodulator, so this is perfectly acceptable. In the context of test and measurement applications however, particularly monitoring applications, the ideal test receiver will measure all PLPs and Layers. It is important therefore to know the capabilities of test and measurement equipment with respect to monitoring multiple PLPs or Layers; test equipment built around commercially available demodulators may only be capable of performing some measurements in "round robin" fashion, i.e. looping through PLPs or Layers of interest periodically.

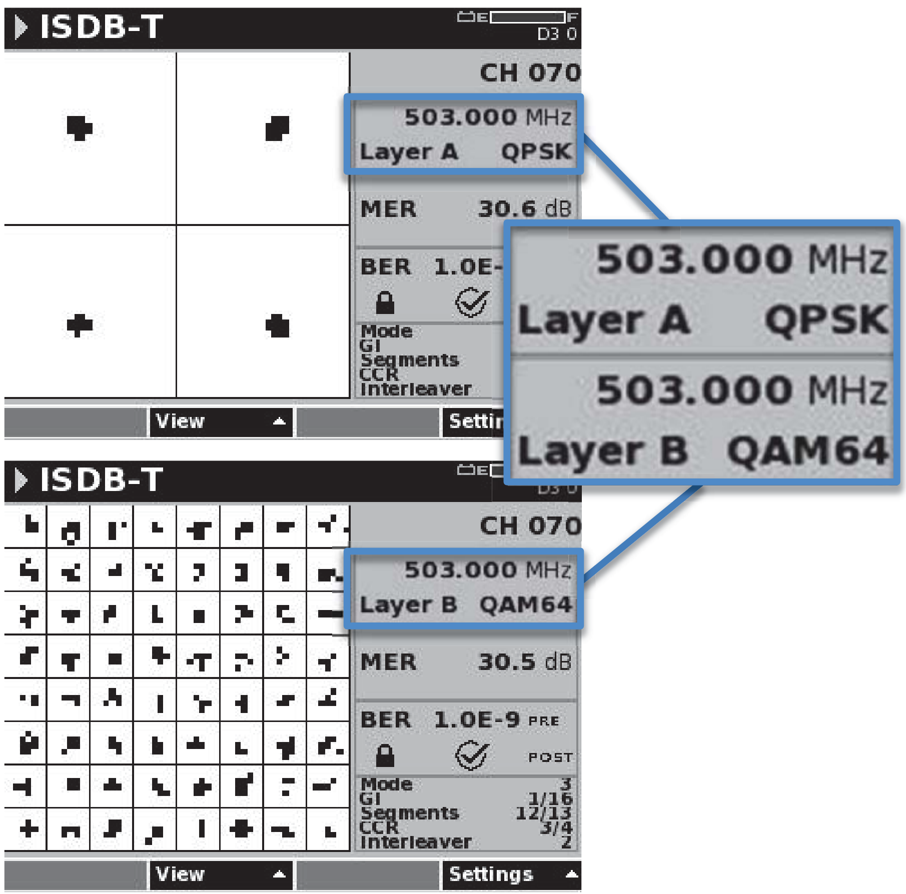

Figure 5 – When hierarchical modulation is employed on a given channel—multiple Layers in the case of ISDB-T as shown here—the constellation and BER measurements are made based on selection of the Layer to be tested [8].

Another impact of OFDM modulation when combined with advanced FEC is that whereas traditionally two BERs have been measured by test equipment, with OFDM utilizing LDPC+BCH error correction there will actually be three error rates: the error rate before any correction ("Pre-FEC") the error rate after LDPC ("Pre-BCH") and the final error rate after all correction ("Post-FEC"). The Post-FEC BER is sometimes referred to as the "uncorrectable errors" and may be shown in some test equipment as a packet error rate.

LDPC error correction is a very powerful error correction technique and is capable of fixing very many errors caused by noisy channels. However, when the input data contains errors, LDPC often leaves a small number of bits uncorrected. The second FEC—BCH in DVB-C2/T2—can be said to be included to "mop up" residual errors following LDPC.

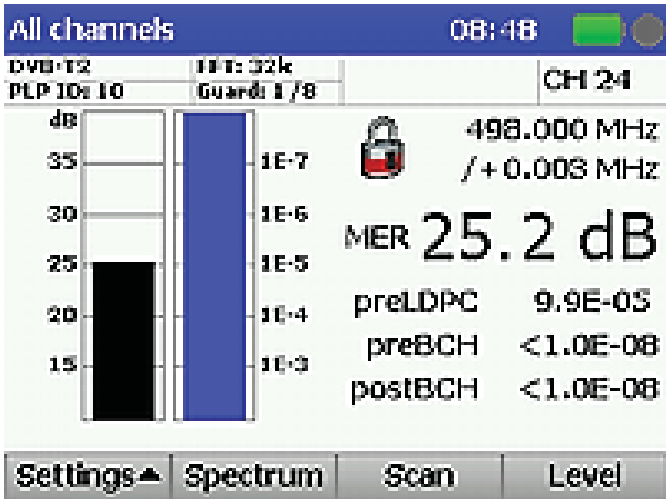

Figure 6 – Advanced, multi-layered FEC algorithms such as LDPC+BCH result in multiple BER measurements per PLP. Here a DVB-T2 test device shows the BER for PLP ID 10 is 9.9E-5 before any error correction, less than 1.0E-8 after LDPC, and also less than 1.0E-8 after LDPC and BCH. DVB-T2 channels with multiple PLPs would have a BER triple for each PLP; the user chooses a PLP to test or tests them all in sequence [5].

Measuring both Pre-FEC BERs is important when BER is to be used as a measure of performance margin because the "cliff effect" characteristic of SC-QAM channels is only exacerbated by the powerful error correction capability of LDPC+BCH. In fact, an SNR difference of only a few tenths of dB separates Virtually Error Free operation from service-impacting degradation [4]. The intermediate BER allows us to see when errors are still occurring after the LDPC correction but at a sufficiently low rate they can be fully corrected by the BCH correction.

Section 5: OFDM constellation

Constellation is a popular measurement for SC-QAM modulated channels because it can help an expert user differentiate between several types of impairments. Modulator-related problems such as I/Q amplitude imbalance, I/Q phase error, and phase noise are each manifest uniquely in the constellation display. Channel-related impairments such as narrowband ingress and additive white Gaussian noise (AWGN) also can be differentiated by an expert constellation measurement user when testing SC-QAM.

OFDM constellation measurements differ from SC-QAM in several important ways. The OFDM waveform contains thousands of subcarriers, each having its own constellation diagram. Displaying these in a single view requires overlaying multiple constellations into a single diagram.

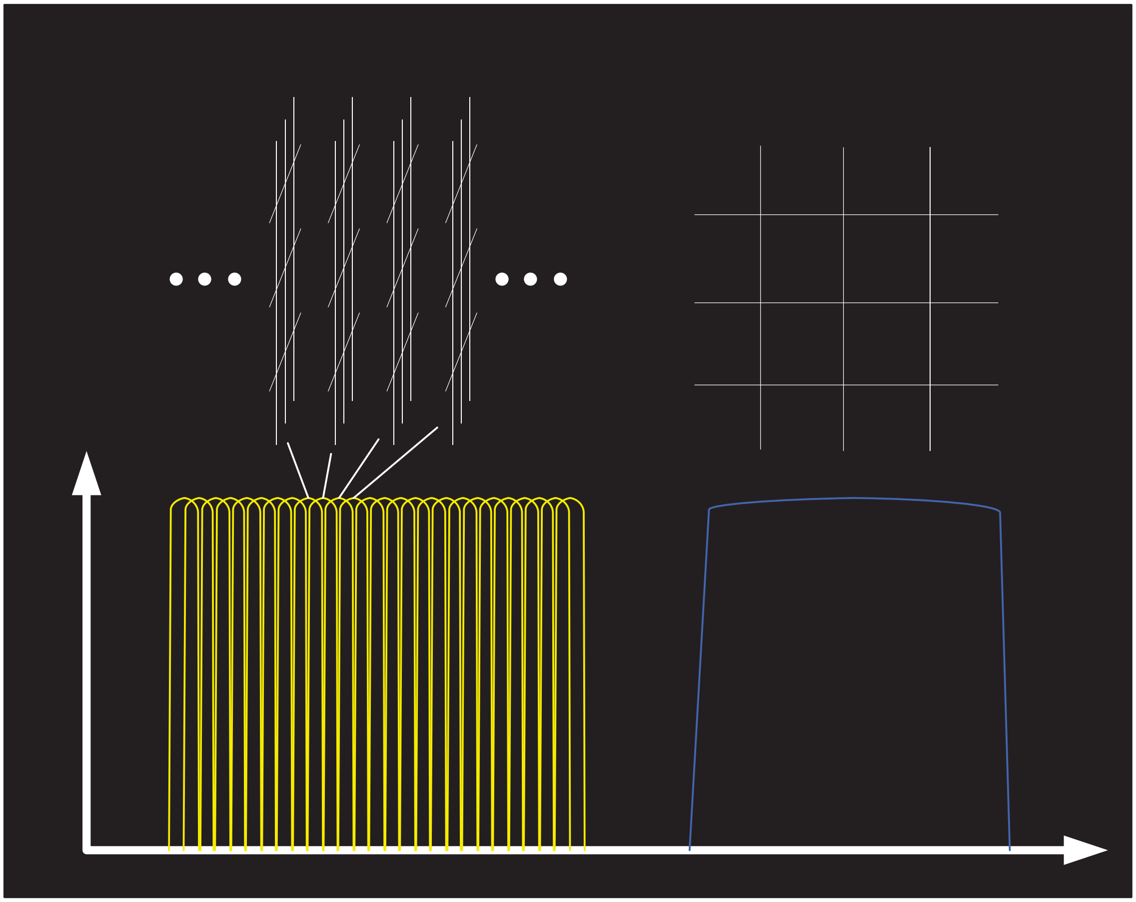

Figure 7 – Unlike SC-QAM modulation, which has a single constellation for each digital carrier (right side) OFDM modulation has a constellation for each subcarrier. Some subcarriers have special interpretation, such as the center subcarrier, continual pilot (CP) and scattered pilot (SP)-bearing subcarriers, and ingress-affected subcarriers.

When a single modulation type is being used across the selected subcarriers, this does not present a great challenge. The main differences from OFDM constellation compared with SC-QAM will be the presence of pilot tones—special unmodulated carriers inserted to the waveform to aide equalization—and the manner in which various impairments manifest. Specifically, modulator impairments such as I/Q gain and phase error appear in the constellation diagram the same as AWGN on all subcarriers except the center subcarrier [4].

Similarly, narrowband ingress will impose the characteristic “donut shape” on only the subset of OFDM subcarriers corresponding to the ingress frequency range. When viewing the composite constellation, AWGN at other subcarriers can mask the ingress and make it more difficult to discern. To view the ingress clearly in the constellation diagram, the user or test instrument must select only the affected carriers for display—no easy feat unless the instrument also provides an easy way to determine which carriers are affected by ingress. Using the constellation measurement to detect the presence and frequency of narrowband ingress, therefore, is of limited utility when measuring OFDM signals.

Combining multiple subcarriers' constellation diagrams into a unified view is straightforward when every subcarrier is utilizing the same modulation format (e.g. 256QAM). When two or more modulation formats are used to carry data for the channel or stream being tested, simply overlaying the diagrams can result in a rather confusing picture.

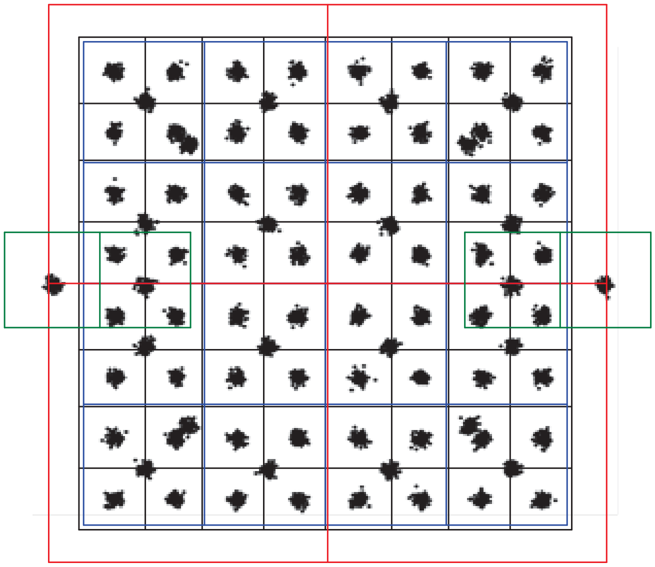

Figure 8 – Plotting all subcarriers' constellations in a simple overlay format can result in a confusing mix of constellations when multiple modulation formats are in use [4]. Here all three layers of an ISDB-T channel are shown in which QPSK, 16-QAM, and 64-QAM are all employed. Pilot symbols' constellations can be seen to the left and right sides of the QAM constellation.

Alternatively, the user can select a single PLP or Layer to be displayed, but again this presumes user knowledge of the PLP/Layer experiencing problems prior to making the selection. In DVB-T2 and ISDB-T, there are typically only two or three modulation types in use at a time, so it is not an unreasonable approach for the user to simply cycle through these or, more often, test the Layer/PLP with the highest order QAM modulation since this will be the Layer/PLP most sensitive to impairments in the channel.

With DVB-C2, many (up to 255) data slices can be bundled to deliver a single stream, and each data slice can contain up to 255 unique PLPs [2]. Granted, using such a large number of PLPs to deliver a single stream is a corner case and not likely to be encountered in real networks. However, it is not difficult to imagine a scenario involving perhaps 8 or 10 Data Slices with one or two PLPs each. In such a network, cycling through all PLPs and Data Slices looking for problems will be time-consuming and of limited utility.

In summary, while the constellation diagram will remain an intuitive view of the quality of the received signal for expert users, it is probably fair to say its utility will be reduced when troubleshooting in OFDM networks compared with SC-QAM. To maximize its usefulness, automatic selection of subcarriers found to be affected by ingress, for example, or those PLPs/Data Slices with the most “interesting” constellation diagrams would be helpful.

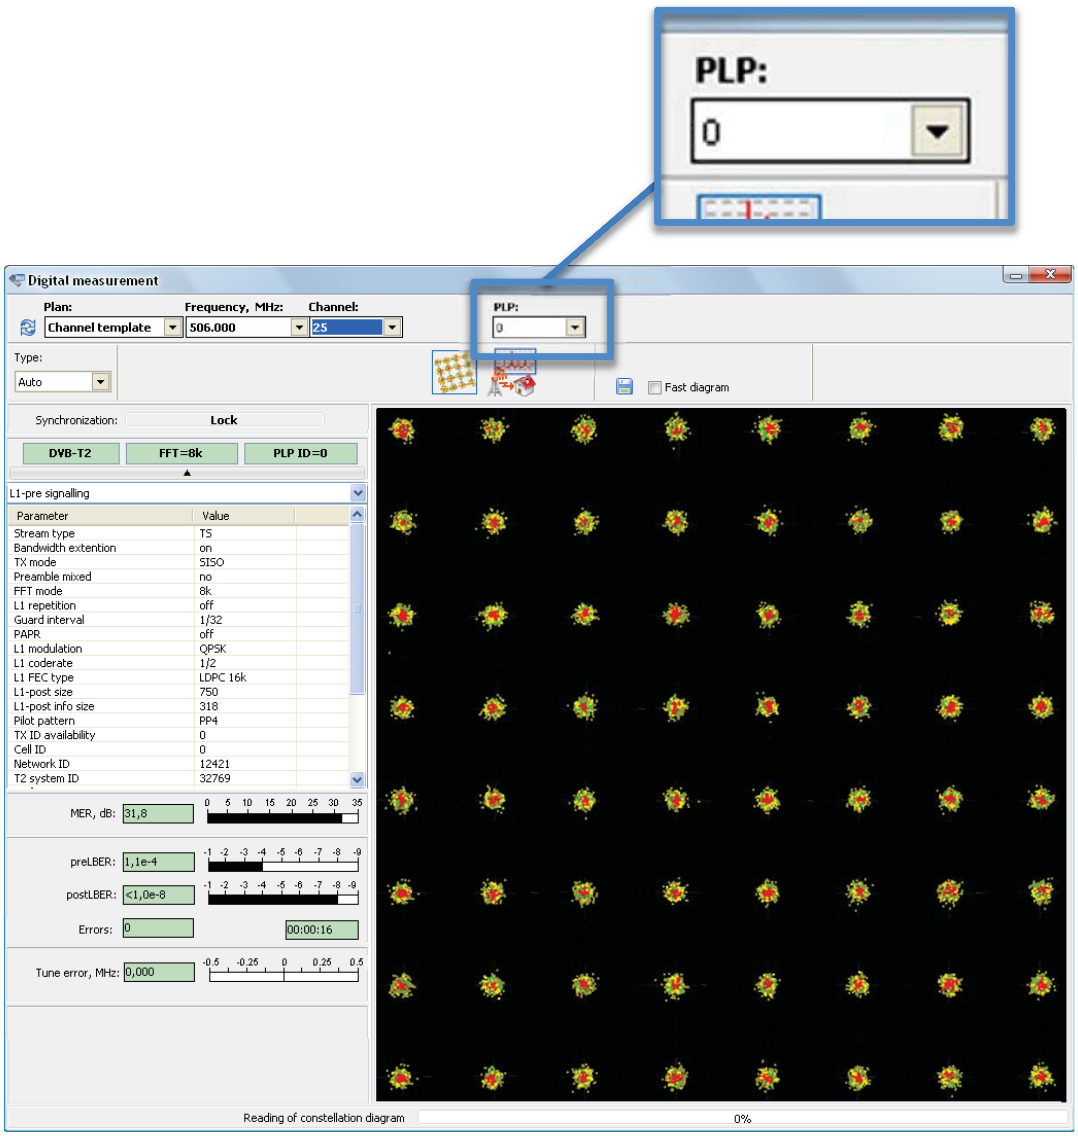

Figure 9 – Alternatively, a specific Layer or PLP can be selected for constellation display. Here the user has selected PLP 0 of a DVB-T2 channel, which in this case happens to be a 64QAM constellation.

Section 6: Equalizer measurements

Equalizer-related measurements represent another category of commonly-used tests to troubleshoot and validate SC-QAM signals. Equalization of OFDM signals differs substantially from SC-QAM equalization, so it should come as no surprise that the use and interpretation of equalizer taps for testing will also be affected.

Unlike SC-QAM, which employs one equalizer with perhaps 16-32 taps, an OFDM receiver has an equalizer for every subcarrier. Within a 6 or 8 MHz span, there are thousands of equalizers, suggesting there will be much more data to be processed for display to the test equipment user. However, offsetting the large number of equalizers is their simplicity: each OFDM subcarrier's equalizer consists of only a single tap: a single magnitude and phase.

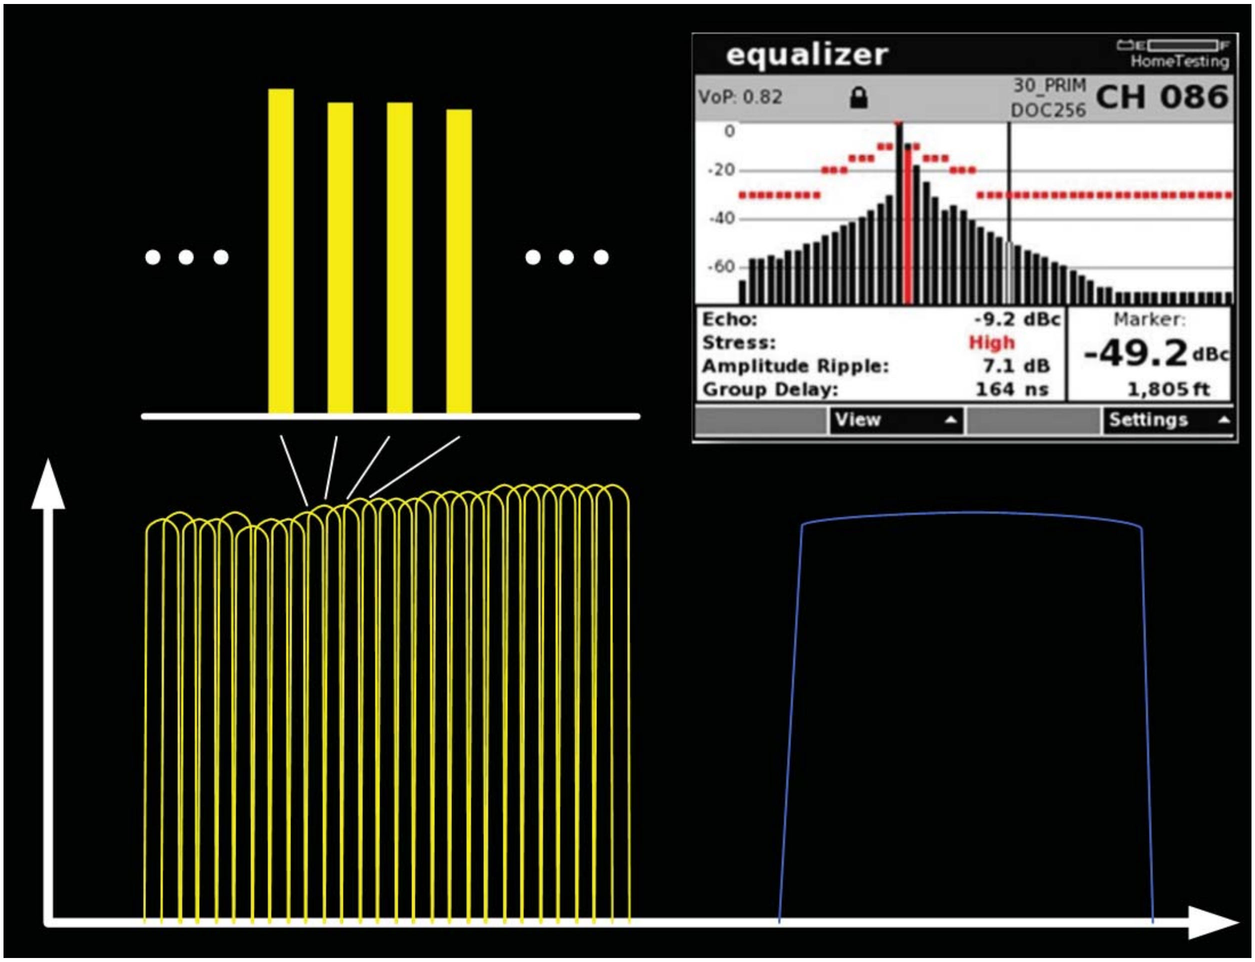

Figure 10 – Whereas the SC-QAM channel (lower right) has a single, multi-tap equalizer (upper right) which natively exists in the time domain, the OFDM waveform (lower left) has one equalizer tap for each subcarrier in the frequency domain (upper left).

How is it possible to equalize an OFDM subcarrier with only a single tap? Simplicity of equalization is a key benefit of the OFDM approach and is possible because the subcarrier spacing is chosen to be narrow with respect to the sharpest frequency domain impairments commonly found in the channel. Remember that each subcarrier is only on the order of kHz in width, by design. It is no coincidence that the frequency response of an HFC network over such a narrow span will be very nearly constant in magnitude response and group delay response.

This is equivalent to the selecting the guard band of the OFDM waveform to exceed the maximum microreflection typical of a channel type, then choosing the symbol rate of the subcarriers such that the guard time is a small percentage of the symbol time. In the case of DVB-C2 for example, it was found that typical cable networks exhibit microreflections mainly less than 2.5 microseconds (approximately 2250 feet) in length [2]. In order to keep the guard time (an overhead) to less than 1% of the total available symbol time, a symbol rate of 450 microseconds was selected (corresponding to carrier spacing of 2.232 kHz).

The consequence of having one frequency domain tap for each OFDM subcarrier is that plotting the magnitude response or phase response of the channel over frequency is trivial, tap values can simply be collected from the demodulator, averaged if desired, and graphed in the test display. Converting the data of an OFDM receiver's equalizers to the time domain however, is not as straightforward. Doing so requires computing the inverse fast Fourier transform (IFFT) of the taps and other mathematical manipulations. The situation is just the opposite of SC-QAM in this respect, where the equalizer taps are naturally in the time domain and must be processed to compute frequency response.

Besides the differences in underlying constellations, the capability of the OFDM equalizer to localize impairments in the time domain—and with a corresponding distance—may also differ from SC-QAM. Annex B 256QAM has a symbol rate of 5.36 Msym/sec; since the tap spacing is normally equal to the symbol duration, 180 nanoseconds in this case, the maximum measurement distance measurable with a 32-tap equalizer is on the order of 32 taps * 200 feet/tap = 6,400 feet in the plant.

In the case of OFDM, the equalizer tap values are determined by the receiver by using the special unmodulated pilot symbols to find the channel response. The OFDM waveform is designed with enough pilots to provide an accurate approximation of the channel. Pilots are not necessary at every subcarrier, only a subset, on the basis that interpolation between the pilots is sufficiently accurate approximation of the channel at those frequencies [2]. The maximum measurable distance corresponds not to the symbol rate, but to the inverse of the pilot-bearing subcarrier spacing.

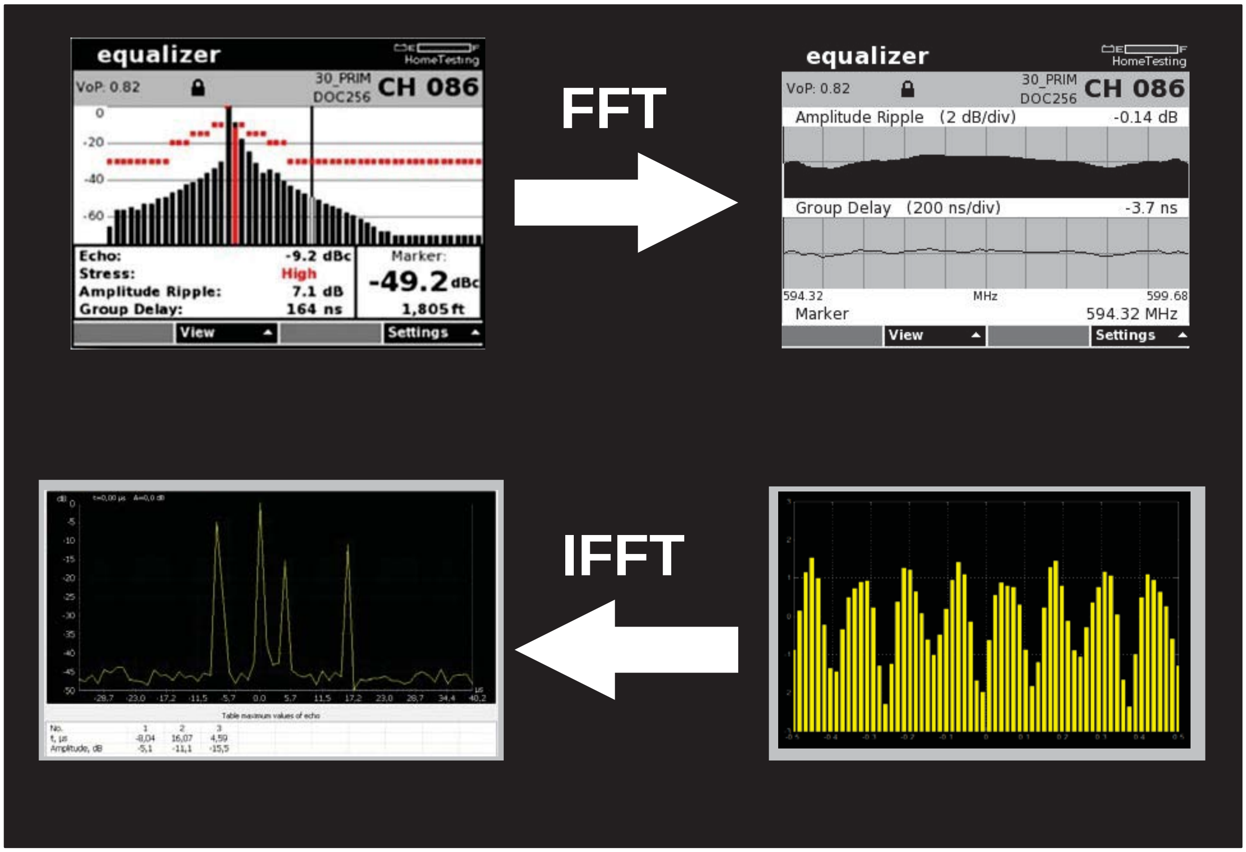

Figure 11 – From its native time-domain format, the SC-QAM equalizer tap values can be transformed to frequency domain results such as Amplitude Ripple and Group Delay via an FFT algorithm and additional mathematical manipulation (upper). The OFDM equalizer works just the opposite; because the tap values exist natively in the frequency domain, the IFFT algorithm must be used to convert them to useful time-domain results (lower).

In DVB-C2 for example, scattered pilots (SP) are inserted at every twelfth or optionally twenty-fourth subcarrier [2], the inverse of which is on the order of 49.8 uSec to 18.7 uSec. This corresponds to a distance resolution on the order of 20,000 to 55,300 feet, a bit farther than the typical SC-QAM equalizer and again, in excess of the total length of the coaxial portion of many HFC networks. Therefore the OFDM equalizer can be at least if not more valuable than the SC-QAM equalizer for determining the distance to remote impairments.

The other important parameter related to using equalizer data to localize plant impairments is the measurement resolution. The highest resolution—smallest separation distance at which two separate microreflections can be resolved—is determined by the tap spacing—the symbol rate—in the case of SC-QAM. The tap spacing, again, corresponds to about 200 feet in the plant. Note that in channels with a single microreflection, especially if post-equalization MER remains high, it is possible to localize an impairment with greater precision than the tap spacing through interpolation. But if multiple microreflections are present nearby one another, measurement resolution becomes a limiting factor.

In the case of an OFDM waveform, the pilot carriers' span determines the measurement resolution. The resolution of a single channel is basically the same as that of SC-QAM, since the pilots span the entire 6 or 8 MHz channel. If however the OFDM waveform is wider band and received coherently, the resolution increases according to the total waveform bandwidth. If eight channels are combined to carry a single waveform, for example, the resolution of the equalizer increases by a factor of eight. If a very wideband OFDM waveform, say 192 MHz is employed, multiple plant impairments less than 3 feet apart can be unambiguously separated.

In summary, the maximum distance measurable by a DVB-C2 OFDM receiver is similar to or perhaps greater than that of a SC-QAM receiver, and the measurement resolution on a single channel is very comparable. This ensures that equalizer tap value processing will remain an important diagnostic tool for the cable network testing. If and when wider band OFDM signals are deployed and received by a single demodulator—say, 192 MHz or wider in bandwidth—the resolution of such a measurement will increase correspondingly, enabling localization of multiple plant impairments at small separation in the plant.

Section 7: Summary and recommendations

The flexibility afforded by the OFDM waveform's use of multiple sub-carriers allows for more efficient use of the coaxial network to deliver high bandwidth services with unprecedented reliability. At the same time, this flexibility introduces complexity for those charged with validating, monitoring, and troubleshooting such services. Potential challenges include more MERs, more BERs, more and potentially dynamically adapting modulation types, potentially confusing constellation diagrams, and new enhanced equalizer-based measurement capabilities.

It may be useful in this context to make a distinction between monitoring applications—used for purposes of identifying and prioritizing problems in the network—and validation/troubleshooting applications in the network and home.

Perhaps the most important implication of the prospect of OFDM-based DOCSIS® networks is the presence of a feedback path from CPE to the CMTS with high resolution physical layer performance data on a regular, potentially low latency basis [9]. This feedback mechanism is implicit to effective use of adaptation for improved reliability.

This data can be combined with information about the type of service being delivered to each endpoint at the time to create a very powerful maintenance prioritization tool. A downstream channel which has been reverted to 256QAM from 4096QAM may or may not represent a service-impacting issue, depending whether the service is a 4K ultra high definition (UHD) video broadcast in a fixed 6 MHz bandwidth or a 10/2 Mbps high speed data service being delivered over 16 bonded channels, 12 of which are running at 4096QAM to maintain a physical layer bandwidth of greater than 1 Gbps [9].

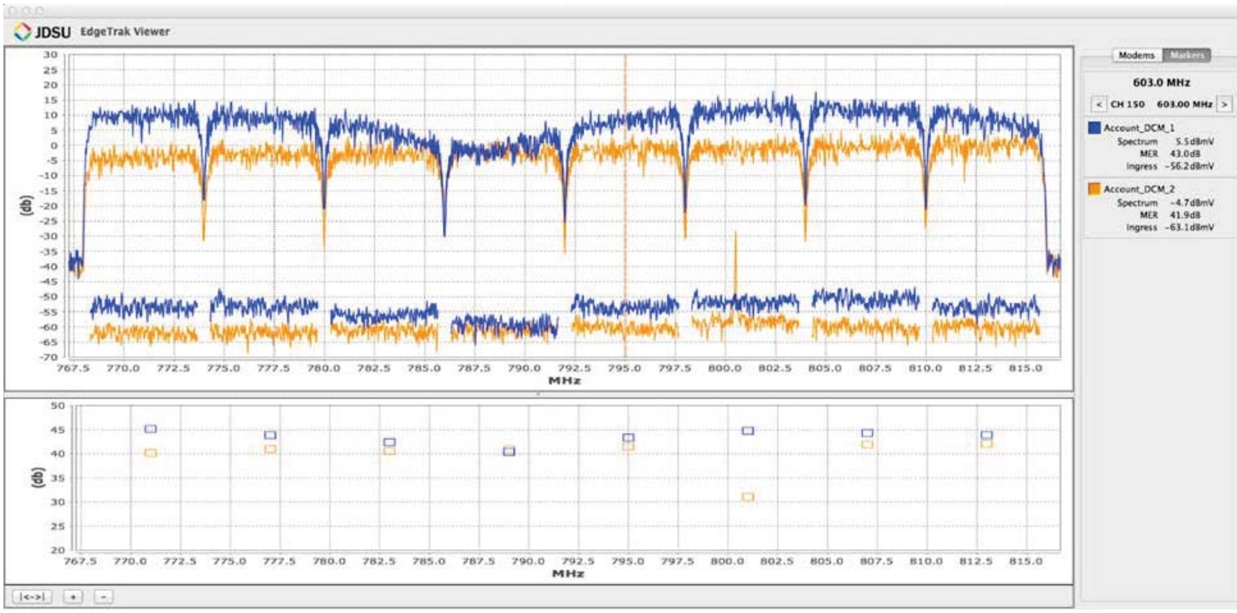

Figure 13 – When downstream receivers' performance data are fed back to the network, to inform adaptive modulation of the subcarriers in the case of adaptive modulation, it can be collected and analyzed to help identify and prioritize impairments in the network or customer premises [10].

The robustness of adaptive modulation provided by OFDM and advanced FEC may provide the luxury of de-prioritizing a service call to a 10/2 high speed data subscriber with only 50 MHz of available bandwidth (perhaps due to atrocious inside wiring) at 15 dB SNR [2]. Using 16QAM it is possible to deliver up to 20 Mbps over a single 6 MHz channel at this MER.

In the context of validating and troubleshooting in the network or customer premises, given this relatively large amount of diagnostic data available from an OFDM receiver and the complex configurability of the waveform, it is more important than ever to ask what data is most critical. At a high level, measurement algorithms will be needed which help simplify this abundance of test data without affecting the underlying measurements' utility.

For example, rather than showing the user the potentially many modulation types in use in a network with adaptive modulation with a corresponding assortment of MERs and BERs to be interpreted, a metric such as bandwidth efficiency may be more meaningful, especially if presented in an intuitive scale. This approach is already applied in some test equipment with application to testing home wiring suitability to support MoCA networks.

Similarly, while Post-FEC BER measurements will always remain the ultimate “catch all” physical layer test, it may be more useful for test equipment to provide some more intuitive measure of the physical layer “margin” with respect to BER, on an intuitive scale. This is much more intuitive than explaining the meaning of “E” and why -8 is better than -7. This too is already being done in some HFC test equipment but will be of increasing value as physical layer complexity increases.

At the end of the day, in the network and at the customer premises, the technician only has control over the quality of the physical medium, namely integrity of the signal path and shielding. If the customer’s issue cannot be traced back to one or both of these root causes, there is precious little that service personnel can accomplish which cannot be accomplished remotely. Whichever new measurements emerge for validating and troubleshooting OFDM signals in HFC, they must be actionable and simple to use.

Bibliography

[1] S. B. Weinstein, P. M. Ebert, "Data transmission by frequency-division multiplexing using the discrete Fourier transform", IEEE Trans. Commun., COM-19(5): 628-634, Oct. 1971.

[2] Digital Video Broadcasting (DVB) Implementation Guidelines for a second generation digital cable transmission system (DVB-C2). ETSI TS 102 991 V1.2.1, Revised June, 2011.

[3] What is OFDM? Ron Hranac, CT Reports. http://www.cablefax.com/cfp/ct/news/ctreports/commentary/What-Is-OFDM_55404.html

[4] Fischer, Walter. Digital Video and Audio Broadcasting Technology: A Practical Engineering Guide. Third Edition, published 2010 by Rohde & Schwarz.

[5] Copper Mountain Technologies IT-15T2 DVB-T/T2 Signal Analyzer Operating Manual. Revision 1.1, May 30, 2013.

[6] The MoCA Quality Index Model. JDSU Application note moca-wp-cab-tm-ae.pdf. http://www.jdsu.com/test

[7] Transmission System for Digital Terrestrial Television Broadcasting. ARIB STD-B31 Version 1.6, Revised November 30, 2005.

[8] DSAM – DOCSIS® 3.0 Datasheet (Japan) DSAM-DOCSIS3MODELS-DS-CAB-TM-JP.pdf, 2011.

[9] DVB-C2: Revolutionising RF Bandwidths' Utilization in Cable. http://www.ict-redesign.edu/fileadmin/documents /1003_SCTE_Broadband_DVB-C2.pdf

[10] EdgeTrak: Remote, down-stream visibility at the true network edge—the CPE. JDSU application note EDGETRAK.PS.CAB.TM.AE, 2013.

Abbreviations and Acronyms

4K UHD	4k Ultra High Definition, video with about 4,000 horizontal pixels
AWGN	Additive White Gaussian Noise
BCH Code	Bose-Hocquenghem-Chaudhuri Code, a class of error correcting codes
BER	Bit Error Rate
CMTS	Cable Modem Termination System
CP	Continual Pilot
CPE	Customer Premises Equipment
CSI	Channel State Information
DOCSIS®	A CableLabs interface specification that enables high-speed Internet services over HFC. The DOCSIS brand for these specifications and devices built to them developed from the specifications' original name, "Data Over Cable Service Interface Specifications."
DTMB	Digital Terrestrial Multimedia Broadcasting, China DTV standard
DVB-C2	Digital Video Broadcast - Cable Version 2
DVB-T	Digital Video Broadcast - Terrestrial
DVB-T2	Digital Video Broadcast - Terrestrial Version 2
FEC	Forward Error Correction
FFT	Fast Fourier Transform
FM Band	Frequency Modulation Band, frequencies of FM radio e.g. 88-105 MHz
HFC	Hybrid Fiber Coax
I/Q	In-phase/Quadrature-phase
IFFT	Inverse Fast Fourier Transform
ISDB-T	Transmission System for Digital Terrestrial Television Broadcasting
kHz	Kilohertz
LPDC	Low Density Parity Check Codes – a class of error correcting codes
MER	Modulation Error Ratio
MHz	Megahertz
MoCA	Multimedia over Coax Alliance
MPEG	Motion Picture Experts Group
OFDM	Description
PLP	Physical Layer Pipe
QAM	Quadrature Amplitude Modulation
RMS	Root Mean Square
RS	Reed Solomon
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SNR	Signal to Noise Ratio
SP	Scattered Pilot