More QAM Channels, Optics, and Bandwidth – Oh My!

Troubleshooting Today's Optically Rich and Expanded Bandwidth Networks

SCTE CABLE-TEC EXPO® 2008 – Philadelphia, PA

Tim Brophy, Ph.D.
Director of Network Architectures
Cisco

Ray Thomas
Principal Engineer
Time Warner Cable

More QAM Channels, Optics, and Bandwidth – Oh My!

Troubleshooting Today's Optically Rich and Expanded Bandwidth Networks

Background

Optical transport for metropolitan, regional, and national networks has used Dense Wavelength Division Multiplexing (DWDM) for many years. DWDM optical transport has become highly sophisticated, achieved amazing data throughput speeds up to 40Gbps, and employed redundancy with lightning-fast failover in case of fiber cuts.

On the other hand, the access part of the HFC network (specifically hub to node) has used 1310 or 1550 nm analog optical laser transmitters for many years. The forward (downstream) path from transmitter to node employs a single wavelength and the return (upstream) path from node to receiver also employs a single wavelength. It is a straightforward architecture and has served operators effectively for years. Optical links utilizing a single wavelength per fiber have served their purpose well and will continue to do so into the future.

During recent years Coarse Wavelength Division Multiplex (CWDM) upstream transmitters have become available for use in nodes as a return transmitter. CWDM upstream transmitters provide the advantage of increasing return traffic without adding fiber and they have been deployed in increasing numbers. In the past two years multi-wavelength (MWL) forward transmitters employing both CWDM and DWDM versions have passed qualification testing and become available for deployment. Progress in laser transmitter technology and optical filter technology has enabled four or more wavelengths in the downstream. Each wavelength can carry a full 1GHz load of analog and digital channels.

The business case for deploying multi-wavelength (MWL) optical systems is compelling. When MWL technology is compared to the costs of constructing additional fiber to increase forward capacity, it is seen that multi-wavelength optics as a "system" is very cost-effective. MWL optics can be used to serve both residential and business customers. It can also free up fiber that is already in use to serve cell tower back haul (CTBH) sites which again postpones the higher costs of fiber construction. Aerial construction costs average around $17,000 per mile in much of the U.S., but reach toward $40,000 per mile in cities. Underground construction costs are generally twice that of aerial with items such as restoration in some areas being a significant factor.

Deployment of multi-wavelength optical systems, however, does mean that the HFC access network will become more complex. This paper will explore those complexities and will present ideas to ease network maintenance and speed troubleshooting.
This added complexity of multi-wavelengths systems is apparent from a diagram. In Figure 1 below [1], the power of a multi-wavelength optical system to add forward and reverse throughput capability is clear. At the same time, the additional components such as multiplexers (MUX) and demultiplexers (DMUX), where wavelengths are mixed, are

potential trouble points. It should be emphasized that while the MUX and DMUX themselves have proven to be extremely reliable components, the signal mixing introduces unknowns.

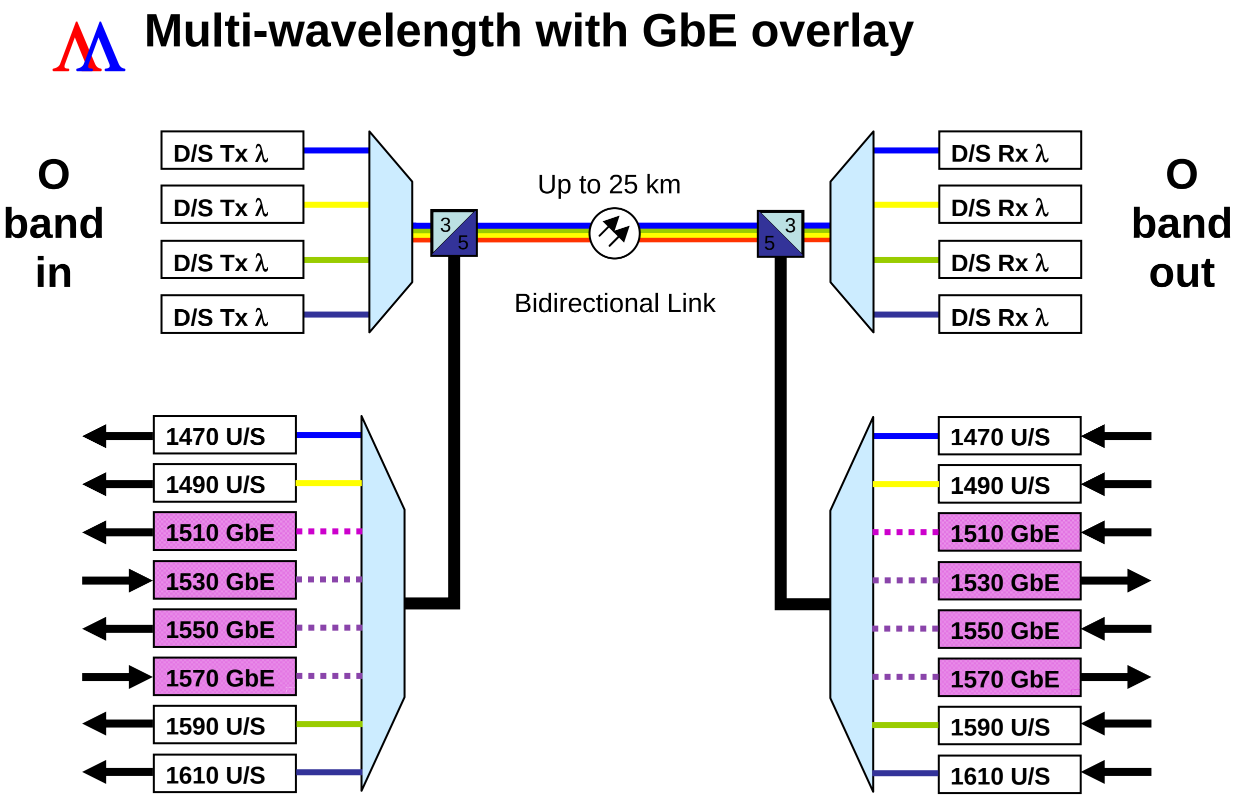

Figure 1. A multi-wavelength architecture using DWDM, CWDM, and GbE [1]

Overview

In this increasingly complex network, there are several key building blocks which, when effectively working together, constitute a framework upon which present-day network monitoring and troubleshooting processes rely. These include network documentation and labeling, and the tools and techniques used to gather and assess pertinent information. As we look ahead, networks will grow to encompass the seemingly disparate characteristics of more uniformity and yet more diversity. These two opposing trends are resolved by recognizing that the diversity of content, application, and consumer interfaces is supported by a distribution infrastructure whose uniformity moves ever-closer to the network edge.

In the progression of network complexity to date, the balance between first installed costs of deployment versus operational expense has resulted in equipment which is lean in the ability to self monitor and correct faults. The trend in transmission products is towards automation of network monitoring and troubleshooting functions, and its implementation is built on those same building blocks of documentation, labeling, tools, and techniques.

The importance of documentation of the network cannot be stated strongly enough. Not only does fiber documentation need to be executed properly during design, construction, turn-up, and launch phases, but it is essential that fiber documentation be kept up to date. Many technicians can tell about their experience of lost or inaccurate documentation costing precious time while troubleshooting an outage.

Another key element is labeling. Labeling enables fast and clear identification of fibers, jumpers, transmitters, multiplexers/demultiplexers, and so on essentially everywhere service personnel will touch the plant. Already important in single wavelength optical links, labeling becomes even more important for links with multiple colors in both directions. Proper labeling promotes fewer mistakes in set up and launch and also helps to avoid inadvertent disconnects when troubleshooting. Labeling goes hand in hand with the documentation steps to aid in rapid reconnection. Labels subject to harsh outdoor environments need to stand up to that environment. A label that has faded due to heat is useless.

The right testing tools will be needed. Test equipment already in use in the access network includes optical power meters, optical time domain reflectometers (OTDR), and portable optical spectrum analyzers (OSA). Handheld optical channel checkers (OCC) have recently become available. An OCC provides some of the features of a more expensive OSA in a lightweight unit that can be used for confirming that each wavelength is present at a given point in the system and that optical power is at the right level. An optical channel monitor (OCM) can be used as a troubleshooting tool or integrated into the distribution network to provide real time active monitoring and control.

Finally, identification of techniques that provide useful information that can help find root causes of problems and potentially prevent outages will be needed. Because signals use common fiber paths and headend or hub shelf hardware (i.e., co-located gear which is highly similar), it will often be impossible to troubleshoot independently an individual link or component. In these cases especially, the necessity of using the combined intelligence gathered from the documentation, labeling, tools, and techniques becomes evident.

Because of the unique challenges it presents, a multi-wavelength system will be used to describe some of the processes used in set-up and troubleshooting of a network. It will rely on and extend many of the familiar optical measurement techniques in use today. Those methods are tried and true when there are individual 1310 or 1550 nm single transmitters, or when the separate wavelengths in a multi-wavelength system act as totally independent sources. In such cases, measured impairments are linear and do not change in the time frame of the measurement itself (although longer term variations such as aging are possible). But they are partly or wholly inadequate when the wavelengths interact with the fiber or each other non-linearly, and the resulting impairment is not proportional to the optical power of the sources involved. For these situations, we need more information. It may be gathered anew, or derived from a more intelligent network.

Until such time as that intelligence is in place, we are often left with only the ability to measure the symptoms of an impairment rather than its cause, and as such, must use our knowledge of what and how impairments may arise to determine what exactly to test in the system. Very often, the documented past performance levels and labeling or screen shots will be the keys to ensuring that the correct fix is not applied to the wrong problem, thereby not resolving—or even worsening—an observed or reported problem. With occasional digressions into considerations that are broadly applicable and consistent with best practices in optical measurements, an extended example of setting up a multiple wavelength system will be used to highlight and evaluate the tools, techniques, expected observations, and future possibilities for troubleshooting networks in the evolving HFC network.

Setting up a system

Multi-wavelength system without optical amplifiers

Referring back to the multi-wavelength system of figure 1, it will be useful to step through how similar systems are installed and brought on-line to ensure adequate performance. Presuming the equipment will be used in conjunction with the vendor's recommended practices, it may suffice to document the origin and intended destination of each of the laser transmitters. Storing birth certificates which contain information regarding optical power, wavelength, test point values, frequency response, serial numbers, etc., is the first step. As status monitoring and network information databases become larger and more sophisticated, much of these data may be available there. Until that automated collection occurs, the initial installation is the optimal time to collect the data manually. Also, although it is a point worth repeating in more detail later, the ability to track changes in initial drive parameters and output values will be very useful. Recognize that not all information collected will be easily available later, and the maxim that "the job's not finished until the paperwork is done" shouldn't be ignored now.

To begin that data collection, it is useful to recognize that even complex systems are in many ways the accumulation of simpler point-to-point systems. So, it makes sense to perform and document the individual wavelength links irrespective of the final network configuration. The shelves with associated mechanical and electrical powering are assumed installed, and discussed here is only the installation and set up necessary for the optical link itself.

Upon installation of the transmitters in the shelf, and proper cleaning of the connectors on the jumpers, optical output power of the individual transmitters is measured and documented. If the capability exists to confirm wavelengths, they should be measured as well. Verify that the output power is compatible with the transmitter model, given the potential for connector loss of ~0.3 dB per mated pair. Excessive loss here may arise from poor connector cleaning—a topic discussed below. Although the point won't be belabored throughout the remainder of the paper, the documentation of measurements should become routine, and the usefulness of those documents cannot be over emphasized.

In the usual case, the individual transmitter's outputs are combined in a wavelength multiplexer. It will be especially important to measure the individual wavelengths separately because in a multi-wavelength system the individual wavelengths will not be available again unless and until they pass through a DMUX (used either for signal distribution or in an optical channel monitor). Again, in the usual case, the optical output at this point will be connected either directly or via a patch panel (whose connector loss and cleanliness should be verified) to the access distribution fiber, and this power level must be compatible with the energy density constraints of fiber non-linearity onset. In the case of a single fiber used for both forward and return signals as shown, the losses through the various ports of this second filter (in both directions) along with composite power measurements should be noted.

Because of its importance in system design, and its potential as a source of confusion in the multi-wavelength power measurements, a digression into some of the underlying causes and means of recognizing non-linear phenomena is warranted [see 4]. Although a detailed determination of the non-linear power level threshold is beyond the scope of this paper, it may typically be assumed that a limit of ~10-11 dBm is compatible with four channel multi-wavelength systems using direct modulation in a link longer than about 15 km. Power threshold levels are, of course, higher for single wavelength systems, in which there is no interaction between or among wavelengths to limit performance. Thresholds here are limited by the single wavelength energy itself. In the case of commonly available directly modulated lasers, the linewidth broadening arising from the chirp of the source in response to the modulating RF current decreases the energy density to values that are low enough to prevent non-linearities arising from *independent* wavelengths. In the presence of very light or no RF loading with low composite modulation depth, or in the case of a very low chirp source, this may not be the case.

External modulation sources are the most familiar example of a very low chirp source, and in this case intentional line broadening is used to suppress stimulated Brillouin scattering (SBS), the lowest energy density dependent optical non-linearity. Similar line-broadening techniques are employed in external cavity lasers (ECL), a source technology introduced recently for FTTH applications using DOCSIS® control in the cable telecommunication's space. A more detailed discussion of transmitter modulation types and their impact is found in [1].

So how does a field tech recognize the onset of non-linear behavior in the system? The evidence of the comparatively abrupt SBS onset above some threshold level can be observed by simple power monitoring at a link receiver as gradually increasing power is launched. Since the SBS reflects light backward toward the (high-power) source the apparent fiber loss—constant for small increments of launch power below the Brillouin threshold—increases with increasing launch power. For changes in input power above the threshold the received power stays approximately constant as the excess power is reflected back towards the source. A simple test to establish the approximate threshold of a long fiber is to measure a link with both low and high power sources. Using the low power launch loss as the actual fiber loss (since launch power is below the scattering onset) the Brillouin threshold may be calculated from the change in measured loss (between high- and low-power launch) subtracted from the high-power launch (figure 2).

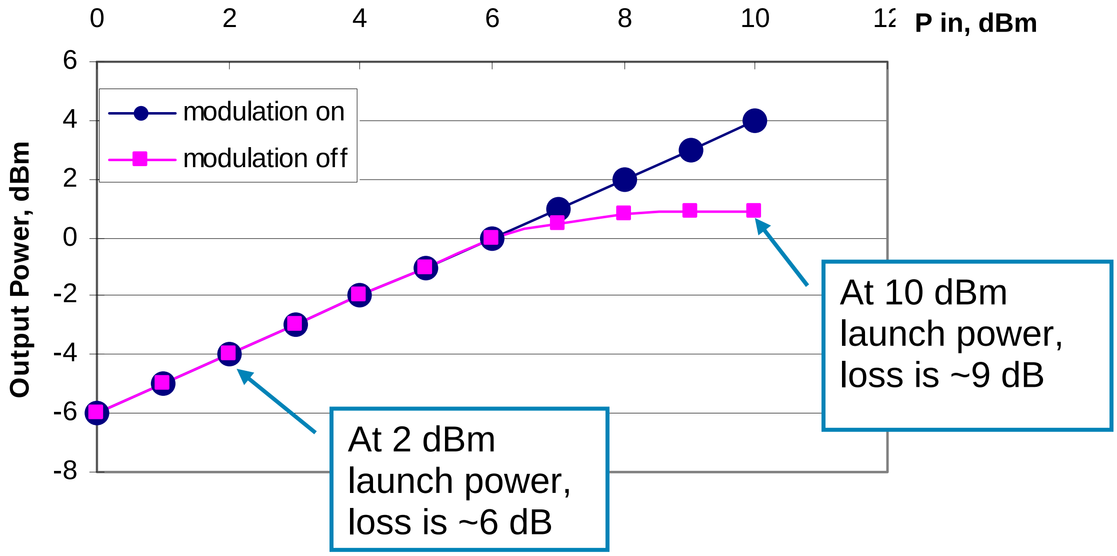

At 10 dBm launch power, loss is ~9 dB

At 2 dBm launch power, loss is ~6 dB

Figure 2. Stimulated Brillouin Scattering (SBS); Since the change in apparent loss is ~3dB, the Brillouin threshold is (10 dBm – 3 dB = 7 dBm). The circles show the RF modulation effectively suppressing the SBS via linewidth broadening.

Whatever the particulars in the various sources, the launch power will be a most useful data point in setting up and troubleshooting networks. Returning to the four wavelength example system, we *independently* measure the multiple sources to ensure compatibility with the non-linear thresholds, and then determine the power uniformity across the channels, which arises primarily from channel loss variation in the multiplexer, supplemented by the connector loss uniformity. Calculation of the composite power and subsequent measurement ensures the (clean) re-connections were made properly. If during setup there is simultaneous access to both ends of a link, it is possible to examine the wavelength-dependent fiber loss separately. Unusual loss at one or more wavelengths, or loss differential greater than that expected, may indicate a discrete reflection, an inadvertent Fabry-Perot cavity (which can lead to signal distortion) or microbending of the fiber itself. While broadband discrete reflections can be verified using an OTDR (see below), wavelength dependent reflections might only be identified by comparing measurements to analytical prediction or to previous measurements.

At the remote end of the fiber link (away from the source), the demultiplexers are often fusion spliced into the link to avoid issues with connectors that may arise due to dirt, other particulate contamination, or temperature variations of the connection alignment which may result in additional excess loss. Again, individual power measurements are made on the separate channels. Values here reflect the combination of launch, MUX, fiber attenuation, and DMUX variations. In a network configuration with separate physical optical nodes used to segment a serving area, fixed location multi-channel demultiplexers may have been supplemented or replaced by single-channel drop filters. While these single-channel filters present some opportunity to measure residual power of other wavelengths, this is more of a component verification than a system troubleshooting tool. Nevertheless, if it is cost-effective to add suitable optical test points at these locations, those test ports could be connected to test gear when available.

Systems using an optical amplifier

While the recent deployments of multiple downstream, full-band wavelengths has presented new challenges, many existing systems already take advantage of multiple wavelengths in the 1550 nm band. The commercial development of EDFA technology combined with maturing externally modulated transmitters led to system designs that take advantage of DWDM components. The additional passive loss and the use of multiple wavelengths feeding a single receiver were mitigated in part by a more widespread use of optical amplifiers, and their use presents unique circumstances for troubleshooting and monitoring a link. This brief detour explores the impact of multiple versus single wavelength use in fiber amplifiers.

In a link that uses an optical amplifier, measurement must take the performance characteristics of the amplifier into account. It is illustrative to use a multi-wavelength system with an EDFA as an example. EDFAs use a gain medium (the erbium doped fiber) which is energized by a laser whose power is able to be transferred to a signal propagating in the same fiber. All EDFAs have a minimum signal input power below which the amplifier’s energy is distributed over the entire possible gain spectrum rather than to the (too low power) desired signal. There is typically an even higher power required to ensure the amplified signal is not degraded by noise. Those power levels are the composite power into the amplifier and not that of each channel, so it is that total power that should be measured. But EDFAs are also characterized by their gain and their saturated output power (the maximum power able to be delivered), which affect individual channel performance. In an incrementally-evolving multi-wavelength system, in which one or more wavelengths are added as capacity demands, the amplifiers are generally operated in constant gain mode, in which each wavelength experiences the same gain. But each channel gets its full gain only if the total output power required does not exceed the capability of the amplifier.

By contrast, EDFAs which will amplify only a single or fixed number of wavelengths are operated in a constant power mode, in order to take maximum advantage of the output power capability. If more channels are added, the total power remains constant, but the power (and gain) per channel drops. An example will be useful to show the impact. Assuming equal power in each of the wavelengths, a WDM system operated in constant gain mode will require 16 dB of additional output power to provide the same gain to 40 wavelengths as it provides when only one wavelength is present (16 = 10*log (40)).

Total Power = Power per channel + 10 log (number of channels)

So if the gain is 18 dB and the saturated power is 20 dBm, the maximum input power of a *single* channel is the difference in the cumulative input power (output power minus the gain) less the channel headroom (see figure 3).

$$\mathbf{P_{ch}\ input,\ max = P_{sat}\ output - gain - headroom\ for\ channel\ count}$$

Numerically,

$$\mathbf{P_{ch}\ input,\ max = 20\ dBm - 18\ dBG - 10\ log\ (40) = -14dBm}$$

Input powers higher than this demand more than the amplifier can deliver and compress the gain and output power of each wavelength accordingly. If that maximum input power is insufficient to avoid a noise penalty, the options include having the minimum channel count higher (to increase the total input power to the amp), capping the maximum number of channels to be amplified, or using a higher saturated output power device. The last of these improves both the input power per channel and the range of channel counts within the limits of the gain and power of the amplifier, but is not as economical. High power amplifiers are available today, but operational concerns usually dictate multiple outputs to spread that power among the multiple wavelength or power ports. Note that it is often difficult to simultaneously satisfy the minimum per channel input power (for noise performance) and maximum available output power (for channel count loading performance). Power measurements combined with an understanding of optical amplifier operational modes and performance specifications will ensure a smooth initial installation and provide insight into where to look in the event of performance degradation.

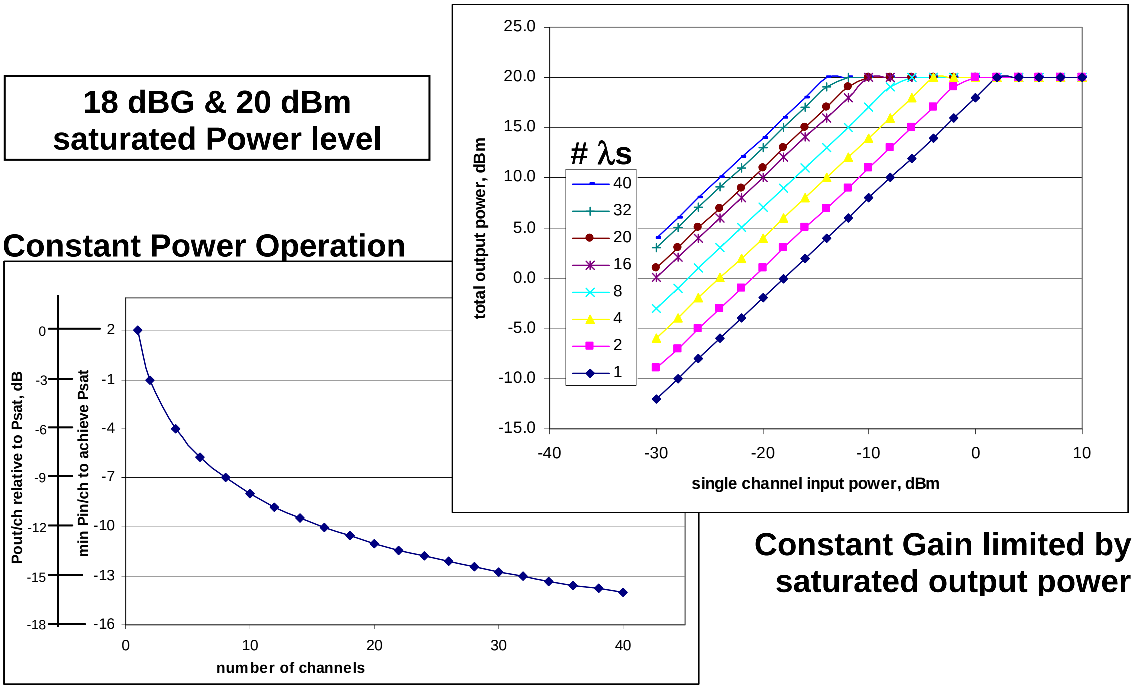

18 dBG & 20 dBm saturated Power level

Constant Power Operation

Constant Gain limited by saturated output power

Figure 3. EDFA modes of operation in a multi-wavelength system; all individual wavelengths are at the same power level.

The above description reveals a significant issue in troubleshooting multi-wavelength networks. It is that information determined by single wavelength point-to-point link-type measurements is often inadequate when multiple wavelengths are present and are not independent of one another. Conversely, knowledge of the aggregate information is often insufficient to reveal a fault that exists on an individual wavelength. Consider, for example, an active power monitor as an alarm trigger. Even in the case of the comparatively benign eight-wavelength system, the measured power decrease with the loss of one channel is less than 0.6 dB (= 10 log (7/8)), which may be on par with expected variations and measurement tolerances, thus rendering an alarm on this basis nearly useless. Clearly, different approaches are required in different network configurations; matching the approach to a particular network requires detailed knowledge of the link.

But even these variations, as difficult to distinguish as they may be, are at least linearly related to the various signal powers. In the case of the multi-wavelength system of figure 1, the non-linear interaction between and among the wavelengths can lead to performance limitations that are at least partly independent of the degraded signal. It is one of the primary motivators for such extensive individual wavelength measurement and documentation recommended. Details about how the interaction of the multiple wavelengths with the fiber characteristics and with each other and how that in turn leads to measurable performance limitations can be found elsewhere [4]. Here it is sufficient to note that those interactions can lead to limitations that are most readily seen and measured in the RF domain, and is the ultimate performance criteria.

As related earlier in the digression on SBS, the line broadening of directly modulated transmitters prevents non-linearities only from *independent* sources. Here, the fiber link itself provides the mechanism whereby the individual wavelengths interact and (non-linearly) affect one another. Specifically, phenomena such as Raman crosstalk and four wave mixing (FWM) lead to CSO limitations and CNR degradations, effects that may first be reported as a degraded picture or tiling service call. These will be verified by traditional RF techniques, but tracking down the root cause may be more problematic.

As an example, consider a four wavelength system with signal wavelengths operating close to the zero dispersion point (ZDP) in the fiber. Multiple wavelengths in this range will generate FWM products of high amplitude and whose locations depend on the originating wavelengths. Recall that FWM is, in one sense, an optical equivalent to CTB. There are some products generated at frequencies $2f_1 - f_2$ (these are the optical frequencies and are reciprocally related to the wavelengths). Because this FWM product is doubly dependent on the f_1 frequency, it will move twice as far (in wavelength) as wavelength 1 itself moves. The difficulty is that this change in wavelength 1 (a shift perhaps due to aging of the laser) may have no adverse impact on wavelength 1, but the FWM product which has moved closer to wavelength 2 has degraded the CNR there (figure 4).

If the wavelength 2 link suffers a degradation, it is natural (in a single wavelength system) to look for the cause in that link alone. But, if only wavelength 2 is examined, there is as yet no simple way to identify this source of degradation with certainty. It may be helpful to compare the operating and fundamental parameters of the lasers sharing the link to the data collected at the time of install. In that way at least, the output power and drive conditions of wavelength 2 could be verified as unchanged; however, for wavelength 1, the shift may have been accompanied by an associated increase in bias current to maintain an output power and which had the secondary effect of shifting the wavelength. It is clues of this level of seeming obscurity that will help determine the root cause.

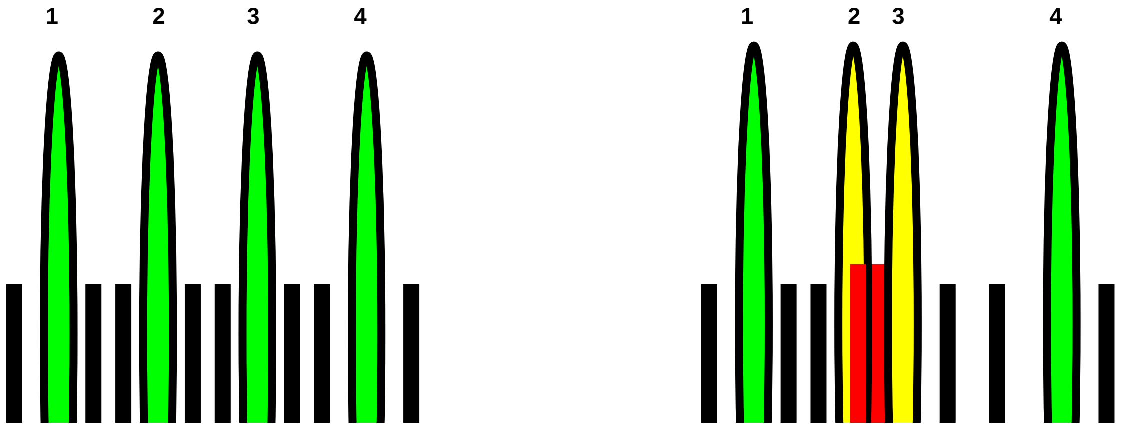

Normal spacing of DWDM wavelengths and beat products.

Lambda 3 has drifted closer to lambda 2. The spacing and amplitude of optical beat products moved with L3 and now interfere with L2. Although L3 was the wavelength that moved, it is still received normally. Although L2 transmitter continued to operate normally and did not drift, the L2 link was degraded significantly by the L3 beats. It would be easy for a technician to mistakenly identify L2 as the bad optical transmitter when in fact it was L3 that degraded.

Figure 4. Example showing a stable link being degraded by a different transmitter

This inability to distinguish with certainty the source of an impairment leads to another question which has both immediate and long term implications, especially in WDM systems. It is the issue of the best method to provide spares for transmission gear. While it is simple enough to replace a defective or questionable transmitter in a single wavelength point-to-point system, this is not so for multi-wavelength networks. As is clear from the example above, the transmitter exhibiting the performance issue may not even be the one that degraded the link, and the reduced performance may be the result of system-level interactions. The options include using a purposely reserved wavelength with a dedicated port on the MUX/DMUX, or choose an “off-grid” wavelength and the upgrade port on the passives. The former option is quite practical in many-channel systems such as C band DWDM when total capacity is not an issue. Then, designating one channel to be used only as a spare enables field techs to carry only that transmitter with them.

Configuration and re-routing the dedicated spare wavelength is also straightforward but a bit time-consuming. As mentioned earlier, the node-end of a link is often fusion spliced and in a separate splice case. This solution would require access to both ends of the link, and an intentional cleave of the unbroken fiber and subsequent splice of a new fiber pigtail to the (working) receiver. Further, if the original wavelength plan is to be restored for operational consistency, a second trip would be required to swap out the spare at a future time. Sparing for the downstream analog multi-wavelength solutions by using the upgrade port with a non-interfering fifth wavelength suffers the same constraints, and may be made even more difficult given the aging characteristics and other non-linear interaction effects with the fiber. In this case, the most rapid and surest plan spares one of each of the four wavelengths, for which there is both guaranteed interaction limits and benefits from rapid implementation with access required only at the source end. CWDM and DWDM return links can benefit from either of the two sparing plans, since a truck roll to the node is unavoidable.

Some Tools of the Trade

In the setup example above, the optical power meter indisputably demonstrates its usefulness as a tool with which every field technician troubleshooting the optical portion of a network must be intimately familiar. In a similar category, but somewhat more restricted based on the higher cost, are portable optical spectrum analyzers, with their handheld versions, the optical channel checker. What follows is a description of some additional tools which prove valuable in installing and troubleshooting optical links.

Optical Connector Video Microscopes

One area of optical measurements that should be emphasized is the cleanliness of optical connectors. This was already important in a single wavelength world. It becomes even more critical with the deployment of multiple wavelengths in the access network. Each and every connector should be inspected for contamination or other impairments of the end face before connecting. Note that we did not say "clean before connecting." If inspection shows that the connector is already clean, swiping it from force of habit is not necessary and could even be detrimental.

Significant progress has been made on development of optical connector video microscopes. One company makes a very small portable unit that plugs into a notebook computer via USB port. The images in Figures 5 and 6 were taken using this small video microscope. The accompanying software provides automated estimation of the level of contamination of the end face by analyzing three zones. Figure 6 shows the results of a badly contaminated end face that clearly delineates the three zones—analysis of an end face that was badly contaminated to clearly demonstrate the three zones. The images can be stored electronically and used for future reference. One may, for example, compare current images of connectors with images a year old if contamination from a dirty environment is suspected.

Fiber connectors used for analog cable telecommunications signals are almost universally of the angled physical contact (APC) variety, such as SC/APC. Use of angled connectors ensures an excellent optical return loss even if the connectors are not terminated. For the "flat" connector styles, such as UPC (Ultra-polished physical contact) or SPC (super-polished physical contact), which have excellent return losses when new and connected to an additional fiber, the reflected power drops to only 14 dB below the incident power if the connector is open, that is, exposed to air rather than another connector. Moreover, the fact that the core of the fiber is perpendicular to the direction the light is traveling means that this reflected light is more readily re-captured by the core and is able to propagate with low loss in an unintended direction, back toward sensitive sources or into the gain section of amplifiers. Angling the connector end face suppresses the capability to re-capture the light to maintain the high return loss required.

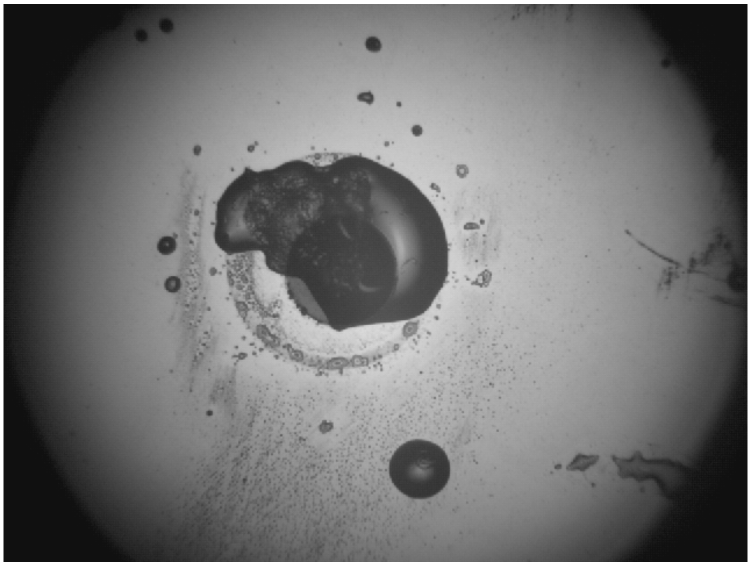

Figure 5. Contaminated end face

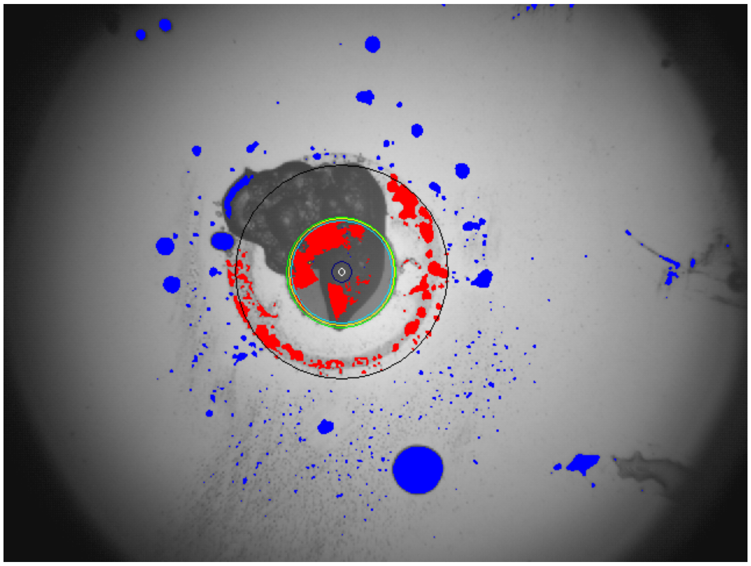

Figure 6. Contaminated end face with zone analysis

In a system with a flat connector, the cleanliness issue can lead to more than just poor return loss. If a small particle of dirt keeps the end faces of two connectors separated, the gap between the ends becomes a Fabry-Perot (FP) cavity. In lasers, this cavity is a useful means to contain light in a gain region to build up power, but in transmission fiber it can cause signal degradation. The throughput of the cavity is wavelength dependent, so a multi-wavelength system may be particularly degraded, but even for an impossibly narrow wavelength source, multiple reflections can create a time-delayed version of the desired signal. This time-delayed signal is seen as CSO in analog video signals. Because the size of the cavity may change with temperature or other stresses, the symptoms themselves may be time-dependent, making their observation and location difficult to detect. Distortions like this are not unprecedented—in the 1990s, a time-varying CSO was caused by polarization mode dispersion, a time-dependent delay due to polarization.

Optical Time Domain Reflectometer (OTDR)

Optical Time Domain Reflectometry in a link is essentially a power measurement of light reflected by a fiber back towards a source. The reflections are small and may occur from a discrete discontinuity or from distributed reflections in the glass itself. These distributed reflections in a long fiber create a background level which can mask weaker signals. The primary sources of discrete reflections in the fiber are splices, connectors, or other discontinuities. Fiber attenuation appears as a slope in the power of the back-reflected signal, consistent with the increasing length of glass the reflected light traverses. The OTDR is most useful for looking at links before they are connected because, if used in live plants, the launched pulses of light can interfere with video or data traffic. This limits the OTDR's use as a troubleshooting tool. Further, in a multi-wavelength environment, the limited number of source wavelengths for the OTDR prevents end-to-end paths from being measured. OTDR traces stored before a link was activated may be useful for comparison if troubleshooting the link at a future date.

OCM: the brave new world of continuous monitoring

Optical channel monitors (OCM) have been employed only rarely in cable telecommunications networks to date. In their simplest configuration, these units connect to a transmission fiber and will measure the power level of wavelengths traversing it. More sophisticated units can measure the optical signal to noise ratio (OSNR) of each of the wavelengths, correct the power measurement for a temperature and wavelength, and track changes in the power and/or wavelength over time. Units are installed as a permanent part of a physical network deployment or as a portable troubleshooting tool, which would connect to a network via an optical tap (used as an optical test point) located at strategic locations throughout a network.

A common physical incorporation of the OCM in digital data networks is to monitor traffic at an add/drop node. Four separate connections and a shared control monitor eastbound and westbound input and output, thereby allowing determination of the dropped and added wavelengths, their OSNR, and their power. Supplemented with appropriate software, units such as these are able to feed a database that would allow fundamental optical operating powers to be tracked over time, enabling proactive alarming if a potential fault is suspected. Thus far, however, the cost of the hardware, and the lack of software integrated with currently supported management systems has prevented more widespread deployments.

To remain in step with the increases in traffic volume and the complexity of delivery schemes, automation in the network management of HFC systems must also become more sophisticated. In such an automated scenario, a cost-effective OCM will become more frequently used if not ubiquitous. System-design link budgets which include optical test points distributed at critical locations throughout the network are not yet common; the drive for cost-effective infrastructure dominates the continuous monitoring capability implied by the use of OCM technologies.

Summary and Indications for Future Direction

Troubleshooting in today's optically-rich and expanding bandwidth networks is a topic both diverse and deep. Virtually every part of the network is evolving at an accelerated pace and every change, no matter how seemingly minor, has a ripple effect that snowballs into a new round of re-evaluation. The physical infrastructure has the double burden of long depreciation period and underlying network support, meaning it needs to work for a long time without failures. This paper highlighted a few of the many perspectives that could be covered.

While the physical layer itself is incorporating new technologies for the delivery of increasingly individualized bandwidth, the tools in the hands of the technicians who keep it running have barely kept pace. One focus was on identifying the importance of documentation and labeling, cleanliness and measurement, which takes on added significance in a multi-wavelength environment in which direct and independent access to the individual carriers is not consistently available. Opportunities to use the information collected as an aid in systematically identifying faults and their cause was highlighted.

Particular attention was paid to networks using multiple wavelengths, as they present the greatest complexity and difficulty in identifying where problems originate. The backdrop of the multi-wavelength system was used to identify techniques used in installation that benefit the installation itself and troubleshooting later. As details not regularly encountered, some identifying information regarding non-linear phenomena and high power effects were presented, as well as considerations for the evolution of a network from a few to 40 wavelengths.

Finally, an example demonstrating the potential to erroneously identify a working link as defective was presented, and recommendations for measurements to correctly identify the real cause were suggested. This led to a discussion of tools currently at the disposal of the workforce, and some that are just being introduced. Ultimately, the desire to provide a network which is not only more robust, but which will monitor, report, and correct for flaws automatically, extended a discussion on tools into the potential for how future networks might be monitored.

The realization of those future networks is more than a question of timing. Although some tools for the physical network are available, the expert systems for fault identification and network- and traffic-aware intelligent management systems are either unavailable or not yet cost-effective. Some of the cost issues stem from the range of variables that need to be tracked, and some from the timing over which to track them. At a minimum, software that proactively monitors the system will need to perform several functions to be useful to multiple system operators:

1. automate the manual collection of data that are critical and informative indicators of network health,
2. store, retrieve, and correlate data for analysis of active (instantaneous) health, aging, usage trends, traffic-routing, and utilization,
3. use data as a predictive tool to trigger alarms in the case of “high-failure” probability tendencies on both the component and system level, and
4. potentially pro-actively re-rout links based on the results.

The second of these functions could be thought of in terms of the time frames that network information is to be used as millisecond, daily, seasonally, annually, and over equipment lifetimes. The capability to pull component or location specific information out of system level data suffers from the same difficulty as tracking a ceiling drip back to a roof leak does, namely, that symptoms can be misleading.

Until these future monitoring and control systems move closer to realization, HFC plants will continue to be evaluated by element managers with a mixed record of data collection, management, and forecasting. Putting the information from the element managers into the hands of well-trained field technicians is presently the best way we have of ensuring our complex systems stay healthy and operational. Improvements to the tools we put into their hands, while useful, is no substitute for training, and for the ability to extract an informed judgment about what *could* cause some particular observed symptom. As networks evolve in complexity and size, so too must the physical, computational, optical, and intellectual toolkit at the disposal of modern cable technicians.

List of Acronyms

APC – Angled Physical Contact
CNR – Carrier-to-Noise Ratio
CSO – Composite Second Order
CTB – Composite Triple Beat
CTBH – Cell Tower Back Haul
CWDM – Coarse Wavelength Division Multiplexing
DMUX – Demultiplexer
DOCSIS – Date over Cable Service Interface Specification
DWDM – Dense Wavelength Division Multiplexing
ECL – External Cavity Laser
EDFA – Erbium Doped Fiber Amplifier
FP – Fabry Perot
FTTH – Fiber to the Home
FWM – Four Wave Mixing
GbE – Gigabit Ethernet
HFC – Hybrid Fiber Coax
MUX – Multiplexer
MWL – Multi-Wavelength
OCC – Optical Channel Checker
OCM – Optical Channel Monitor
OSA – Optical Spectrum Analyzer
OSNR – Optical Signal-to-Noise Ratio
OTDR – Optical Time Domain Reflectometer
QAM – Quadrature Amplitude Modulation
SBS – Stimulated Brillouin Scattering
SC/APC – Subscriber Connector/Angled Physical Contact
SPC – Super-polished Physical Contact
UPC – Ultra-polished Physical Contact
USB – Universal Serial Bus
ZDP – Zero Dispersion Point

References and Acknowledgements

1. Not all WDM systems are created equally, Timothy J. Brophy and Fernando Villarruel, SCTE Cable-Tec EXPO, Philadelphia, Pa, June 2008.

2. Multi-Wavelength Access Networks: A Practical Guide to Implementation, Ray Thomas, Venkatesh Mutalik, National Cable Telecommunications (NCTA) Technical Papers, 2007

3. Introducing LcWDM – The Next WDM Technology for the Cable Industry, Oleh Sniezko, Sudhesh Mysore, Charles Barker, Aurora Networks, National Cable Telecommunications (NCTA) Technical Papers, 2007

4. The Delicate Balance, Tim Brophy, Fernando Villarruel, Kuang-Yi Wu, online at http://www.scientificatlanta.com/multiwavelength_paper, Cisco White Paper, 2007, and abridged as "*A Delicate Balance,*" Communications Technology, Vol. 25, No. 1 , p. 33, January 2008.

5. Private communication, Wim Mostert, Aurora Networks, 2005-2007

6. Gigabit IP and Analog/QAM Traffic Compatibility in HFC Networks - A case for Modern physical layer convergence, Shamim Akhtar, Doug Weiss, Comcast Corp., Venk Mutalik, Zoran Maricevic, Marcel Schemmann, David Heisler, Don Wesson, Liyan Zhang, C-COR Inc., SCTE Conference on Emerging Technologies, 2006

7. Cable Television Technology, Second Edition, Chapter 13, Wavelength-Division Multiplexing, The Morgan Kaufmann Series in Networking 2004, by Walter Ciciora, James Farmer, David Large, Michael Adams

8. Enhancing cable architectures with WDM - Grafting WDM onto existing cable systems, Venkatesh G. Mutalik, Communication Engineering and Design (CED Magazine), February 1997, Page 54, WDM Series

9. ITU-T Recommendation G.695 - Optical interfaces for coarse wavelength division multiplexing applications: Series G Transmission systems and media, digital systems and networks (01/2005)