PLUGGABLE OPTICS FOR HFC

Fernando Villarruel, Fred Hirt, and Kuang-Yi Wu
Cisco

Abstract

Our time is that of an insatiable appetite for fresh, new information, taking and giving, alone and in unison; from grandma who Skypes with the grandkids, to the boy that has never known linear television, to the masses who can Tweet a country to political change in a matter of days, and everyone else on Facebook.

Access platforms are built for their time and the successful ones are a complement of technologies that work towards a particular goal. So what is the access platform of our time? What does it do? How does it scale? What key technologies does it integrate to facilitate the services the cable industry aims to provide?

In this paper we propose an optical RF transmitter that allows integration of HFC optics into next generation CMTS platforms facilitating all-IP transmission and allowing power and footprint reductions through the use of 10 Gbps optics, such as is found in XFP packages.

INTRODUCTION

Cable operators are actively creating road-maps toward end-to-end all-Internet protocol (IP) service functionality. Simultaneously, recent technical developments have resulted in various schemes for access delivery via IP video to complement existing Data over Cable Service Interface Specification (DOCSIS). Yet, while there is agreement on IP video delivery as a goal for access transport, there is no broad agreement on the specifics of the access plant hardware required to accommodate IP video delivery *without discarding major portions of the hybrid fiber coax (HFC) network*. Thus, a greatly desired part of any service transition is key technology advances that allow an increase in capacity and performance while maintaining as much of the sunken investment, particularly in the outside plant.

In this paper, as a key part of next generation access platforms, we propose a simple, cost-effective, power and footprint efficient method to enable IP traffic *utilizing pluggable RF optics*. The RF pluggable optics we propose maintain the cable operator's significant investment, from fiber spans, through the node and into the RF plant. This is in contrast to Ethernet-over-coax techniques or schemes to remote the cable modem termination system (CMTS) or RF gateway physical-layer hardware out to the node, all of which significantly impact the nodes and RF plant.

Particularly, the technology advance we propose is integrating the HFC optical transmission into higher-layer platforms, such as Comcast's Converged Multiservice Access Platform (CMAP) or Time Warner Cable's Converged Edge Services Access Router (CESAR) via RF-modulated small-form-factor pluggable optic modules. We focus our attention here on the forward path, but we also advocate integrating the reverse path optics as well, leaving that specific treatment for another opportunity.

PLUGGABLE OPTICS

The concept of pluggable RF optics logically follows from an understanding of the two demarcation points that define HFC access architectures. These are the transition point from baseband digital content transmission to RF modulated transmission

[RF gateway or cable modem termination system (CMTS)] and the transition point at which fiber ends and coaxial cable transmission begins (typically the optical node). We will look these two demarcation points from the perspective of an access network designed for end-to-end IP video transport.

Our discussion includes the architectures expected and technology specific to modulation formats. Within this broader discussion we will answer questions about the viability of our proposition such as,

- Are pluggable RF optics feasible, can they even be built?
- Under what circumstances can they be produced?
- Would other network components be required to change?
- Can pluggable RF optics be fabricated cost-effectively?
- Is a standard attainable, and if so how?
- Are there line power utilization benefits, i.e. is it green?

We propose answers to these and other questions to paint a clear and undisputable picture of pluggable RF optics as the rapid-deployment, cost-effective means of achieving end-to-end IP video delivery.

HFC AND IP

Transmission formats create boundaries and opportunities. In the case of HFC, its strength and flexibility is that it leverages both frequency division multiplexing (FDM) via the RF spectrum and simultaneously time division multiplexing (TDM) via DOCSIS. This unique combination has allowed it to scale from very basic services including broadcast analog transmission, to narrowcast video services, high-speed data, and voice over Internet protocol (VoIP) telephony, without any fundamental changes to methodology of transmission. Now in the wake of an IP services boom, it merits re-evaluating if this FDM/TDM combination is still useful and practical.

In the last few years the capacity of HFC architectures has increased significantly with a migration to large numbers quadrature amplitude modulation (QAM) channels. This migration creates a scenario where the capacity of an all-QAM signal lineup can be competitive with that of any other architecture, even FTTH [1]. Specifically, with a usable data rate of about 38 Mbps per 6 MHz bandwidth 256-QAM channel, the RF spectrum in the forward path as a whole can very easily grow to an aggregate 5.8 Gbps. Nevertheless, it is not only raw capacity, but the simultaneous use of spectrum partitioning and timed availability that multiplies HFC's effectiveness in comparison to other TDM- or FDM-only applications.

A conservative future example is that of 200 homes sharing a full all-QAM forward path spectrum for an all-IP service offering. In this example the leveraging of multiple bonding groups within the RF spectrum and including bandwidth accommodation for reverse path growth via a mid-split segmentation easily allow a very competitive transmission rate of 1 Gbps downstream and 100 Mbps upstream [2]. Most importantly, an IP deployment of this sort has the ultimate benefit that the node transition point and function remains. Thus, when evaluating what to keep and what to leave behind in IP architectures, it is hard to put aside the current combination of RF modulation and DOCSIS.

The previous example and others showing migration to smaller service groups, in conjunction with maintaining RF transmission [3], are in line with optical segmentation techniques deployed during the last few years.

Scaling for future optical wavelengths in service could result in anywhere from four to eight times as many as there are now, leaving on the table the very real and pertinent question of footprint and power availability for projected new hardware. This is a fundamental, practical question that equipment providers must answer as cable operators migrate to all-IP networks.

COMPLEX RF MODULATION (CRM)

In order to achieve maximum bandwidth efficiency in the physical transport layer, high-order (64 through 1024) QAM transport is required. We refer to such *all-QAM signaling as complex RF modulation (CRM)*, as distinct from that of traditional analog video (NTSC, PAL, etc.), quasi-constant-envelope digital signaling (QPSK, O-QPSK, etc.), base-band digital transmission techniques (e.g. OC-192, 10 G, etc.) or mixes thereof. This prevents ambiguities associated with the more general terms "analog" or "digital" transport, which can variously refer to amplitude-modulated vestigial sideband (AM-VSB), various orders of QAM transmission, or the baseband digital format of some digital return path, backbone, metro, and cable's transnational links.

CRM loadings are fundamentally different from the mixed AM-VSB/QAM loads which comprise the majority of current deployments. Intuitively, one expects a uniform CRM loading to be "easier" to transport than mixed analog/QAM content, with CRM yielding more robust signaling, and greater noise tolerance. All of this leading to greater link budget. This is true, but what makes it so?

Access Link Performance Requirements at the Node Output for Existing & Future Access Payload Modulation Schemes		
Performance Parameter [1,2]	**Existing (Analog/QAM)** **• 78 Analog Carriers** **• 75 Carriers, 256-QAM**	**Future (CRM)** **• 153 Carriers of All 256-QAM** [2]
CNR (dBc)	> 50	>40
CSO (dBc)	< 63	< 55
CTB (dBc)	< 63	< 55
MER (dB) [3]	> 37	> 37
BER, Pre-FEC [3]	$< 10^{-9}$	$< 10^{-9}$
BER, Post-FEC [3]	$< 10^{-12}$	$< 10^{-12}$

[1] Analog measurements according to ANSI/SCTE 06 2009, ANSI/SCTE 17 2007
[2] 153 QAM carriers in continuous-wave (CW) mode to measure CNR, CSO, CTB
[3] Equalized QAM, measurements of ANSI/SCTE 121 2006, ITU J.83 Annex B source.

Table 1− Performance Requirements: Existing and Future Access Payloads

Transmission Characteristics for a CRM Payload

To understand we begin by examining fundamental access link performance parameters to see what must be maintained and what can be relaxed when transitioning from analog-rich to analog-free link payloads. Table 1 details performance parameters expected for a current access optical link, as measured at the HFC node. It compares performance for a mixed modulation loading of 75 AM-VSB channels with 75 channels of 256-QAM (representing an existing case for many access networks) to a load consisting of 153 channels of 256-QAM.

The network performance differences between the access loading schemes shown in Table 1 are both subtle and significant. Of particular note are the differences in the carrier-to-noise ratio (CNR), composite second order (CSO), and composite triple beat (CTB) values required for unimpaired transmission in each case. Let us examine the details relating these parameters in order to differentiate between the requirements for existing, mixed analog/QAM HFC access networks and an end-to-end IP video network employing CRM transmission for access.

First, in order to measure "analog" parameters such as CNR/CSO/CTB for a 153 channel 256-QAM CRM load, all QAM modulators used during measurement must be set to continuous wave (CW) operation. Further, such CW level must be calibrated at a level corresponding to modulated carriers yielding a minimum 37 dB equalized modulation error ratio (MER) for the 153 channel QAM load. Although obvious, this is necessary in order to differentiate among linear and nonlinear impairment mechanisms which result in the noise component of the QAM MER. Essentially analog measurements are used to give a more detailed description of the mechanisms responsible for impairing (or limiting) the QAM MER values of an access link.

Second, in general it is understood that analog parameters are necessary, but not sufficient, to yield robust MER values. This is due to the fact that they do not fully account for the effects of phase noise or "quasi-phase noise"-like effects. Thus, there exist instances in which phase noise components greatly determine the QAM MER performance, particularly in the case of very high carrier-to-composite noise (CCN) and CNR values. Such cases occur frequently in mixed analog/ QAM links when large (> -40 dBc) analog distortion products fall near or under a QAM carrier and are resolved by the customer premises equipment's (CPE's) demodulator as non-coherent single frequency components. This quasi-phase noise degrades modulation recovery, thus reducing MER. *This impairment is unique to mixed analog-QAM transmission*, due entirely to the high energy analog carriers producing discrete distortion products. It is important to note that such effects do not exist in all-QAM CRM transmission.

Thus, robust optical links themselves do not contribute noticeable phase noise to QAM signals. Any residual phase and delay impairments beyond the access optical link, due to RF impedance mismatches for example, are adequately compensated for by the QAM receiver's adaptive equalizer.

A third point is that the optics be *approximately* linear for amplitude and phase transmission. This implies no clipping and no compression in transmission along with the avoidance of excessive, variable timing delays. That is, no variable delays on the order of multiple milliseconds, as is the case in route redundancy switching between greatly differing time-of-flight routes. Such changes adversely affect upstream ranging (for CMTS) and downstream latency in VoIP applications.

In the case of clipping and distortion, such issues are routinely dealt with in proper optical transmitter design and calibration. In the case of redundant link delays, they can be accommodated by approximately-matched delays in the redundant link layout. Ultimately, well-designed optical link delays are limited by dispersion, a sub-nanosecond phenomenon which does not contribute significant phase noise to QAM signals at access optics link lengths (sub-100 km).

In addition to the hardware and link considerations, a final point in understanding the performance of CRM signals over HFC access optical links is the nature of the nonlinear components generated. In mixed analog and QAM access transport links, impairments consist of noise, discrete distortion products due to nonlinearities such as CSO and CTB, as well as optical RF crosstalk and beating effects. As previously mentioned, such discrete distortion products can lead to tones lying near or under QAM carriers, which result in degraded MER, while still exhibiting low noise and excellent CNR.

In contrast, a CRM payload's nonlinear impairments manifest themselves as Gaussian noise-like components. That is, the second and third order products are noise-like rather than clusters of composite beats, and *can be considered additions to the noise floor under a QAM carrier*. Further, this principle also extends to multi-wavelength crosstalk components such as optical cross-phase modulation (XPM) and four-wave mixing (FWM) [4].

In summary, a CRM loading offers a new set of choices for the network designer due to the more forgiving nature of the distortion impairments appearing as Gaussian noise-like components. End-to-end IP functionality benefits from both the increase bandwidth efficiency of high order, all-QAM CRM transmission, as well as the relaxed transmission requirements as compared to that of mixed analog/QAM channel loads, shown in Table 1. IP video transport takes maximum advantage of a CRM loading which does not suffer from the out of band discrete distortion beats created by analog channels and exacerbated as drive levels reach non-linear peaking or compression. In mixed loads, those beat clusters falling near or in QAM channels stress decision boundaries and are ultimately problematic for demodulation routines to withstand and correct [5, 6.]

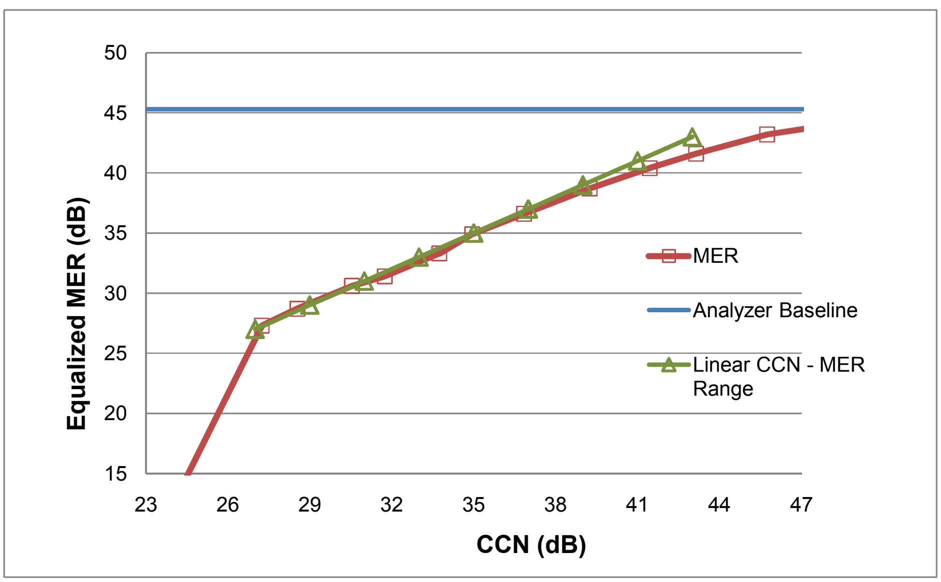

Figure 1 – Equalized MER as a Function of Carrier to Composite Noise (CCN) Ratio

<u>Exploiting the CRM Advantage in DOCSIS Data and IP Video Access Links</u>

The Gaussian noise-like nature of distortion products generated in an all-QAM access link allows the CCN of a CRM loading *to be near-linearly related to the MER*, see Figure 1. In this plot, the CCN and equalized MER are seen to have a linear relationship within the operating range of a Rhode +

Schwarz EFA QAM signal analyzer. The CCN varies linearly form a 28 dB CCN lower analyzer acquisition locking limit (for 256-QAM), to a 42 dB CCN upper limit due to the analyzer's 46 dB maximum MER measurement capability. This observation of a linear relationship between CCN and MER leads to the following two advantages of an end-to-end IP access network based upon all-QAM, CRM transmission.

The first is that the linearity necessary for unimpaired optical transmission is *greatly* decreased when dealing with CRM network payloads, specifically relaxing CSO and CTB for both optical and RF domains, as well as relaxing XPM and FWM requirements in the optical domain. Specifically, CNR is reduced by nearly 10 dB, with CSO and CTB reduced by 8 dB. This opens a number of possibilities for new links and radically different RF transmitter designs, *greatly* simplifying the design, manufacturing and tuning challenges for CRM optical transmitters over those of their mixed-payload analog predecessors. Pluggable RF optics leverages these differences to reduce optical transmitter size and power dissipation while tightening its integration within the headend or hub's content-generation hardware. These reductions greatly reduce shelf space and power dissipation while improving ease of use and sparing issues.

The second benefit is that the linear relationship between CNR/CSO/CTB along with optical cross-talk (XPM and FWM), and MER (recalling that nonlinearities map into CCN which is linearly related to MER) allows the use of the same design approaches traditionally used to make hardware decisions for access hardware links, with the advantage of relaxed noise and distortion goals. That yields positives all the way from design and manufacturing, to deployment and turn-up. From the selection of components used in the transmitters to their manufacturing and tuning, to the link designer, headend tech, node techs and maintenance staff, the CRM transmitter relies on well understood, tested transmitter technology. Despite being end-to-end digital, with DOCSIS data and IP video delivery traffic, networks utilizing 256-QAM designs can target specific CNR, CSO and CTB goals, with the expectation of well-defined QAM MER performance.

One caution regarding the data shown in Table 1 must be raised. It might be implied that if a heavily-analog channel loading does not permit the advantages of a CRM payload, then half as many analog channels would "spilt the difference" and be a good compromise. The reality is that any inclusion of analog channels has the negative effect of producing *discrete* clusters of distortions which can and do peak within the bandwidth of a QAM channel, degrading overall performance. The degradation that occurs from the presence of analog distortion products affects MER in a fashion *which CNR and CCN will not reveal*.

It is necessary that next generation hardware employ all-QAM, not mostly-QAM, transmission in order to fully leverage the benefits which CRM loading brings over mixed analog/QAM loadings. As a practical matter, provision can be made for the addition of two or three service tones, spread throughout the operating bandwidth, without degrading the CRM signal. In summary, CRM payloads allow simplified, straightforward design rules for access links, which can be exploited to greatly improve access transmitters for end-to-end IP video delivery.

<u>Link Considerations</u>

Another benefit of CRM is that it can increase the link budget, the exact amount of which depends upon link parameters such as equivalent Optical Modulation Index (OMI) per CW channel and the desired RF output at the node. Switching to CRM can result in nearly a 3 dB reduction in optical input power to the node receiver since in current networks the QAM channel powers are already 6 dB (RF) lower than the AM-VSB channels. Note that the full 3 dB optical power decrease is not *always* achievable by loading change alone, since the 153 channel QAM load is not being compared to a 153 analog load but rather to a mixed analog/QAM loading.

A further reduction to receiver input power can be made, however, if the QAM loading utilized exhibits a high CCN. In such cases the receiver input can be reduced to the point where the shot and thermal noise components of the receiver dominate at this lower input power. Such a link budget improvement can be used to lower launch power and so lower the non-linear dynamics occurring in the fiber. Since fiber nonlinearities are launch-power-dependent even a 2 dB launch power reduction can yield a significant reduction in crosstalk and four-wave mixing [7].

In cases where the operator has some leeway in accounting for RF power, throughout the RF chain to the home via unused amplification potential or node segmentation, the optical input power into the receiver can be reduced by more than 3 dB, down to -10 dBm or lower depending on several performance factors.

The implication of higher optical link budgets also creates ample space for the reduction of stimulated Brillouin scattering (SBS) suppression, which has been one of the daunting challenges for cable optical transmitters since their inception. SBS is a scattering effect that takes place in fiber when the launch power of a wavelength is approximately greater than 7 dBm. Special and very complicated circuitry has been created to compensate for this issue

IS PLUGGABLE RF OPTICS A REALITY?

So the question remains, can small form factor pluggable transmitters be a part of the HFC landscape moving forward? The answer is a resounding yes. The upshot of reduced linearity requirements for all-QAM channel loads is that it creates a potential tangible shift from the hardware employed to make legacy cable optical transmitters, to making future IP ready transmitters. In particular, there are new opportunities in the mix of components that can be used to reach the desired performance values, in addition to reduced size and power consumption.

For many years now the ability to make cable transmitters has been determined by a few key factors. In the case of directly modulated transmitters (DMTx), their lasers have had to have a minimum necessary linearity and stability dependent both on the growth characteristic and packaging structure, ultimately creating a specific pool of usable lasers and a size threshold for the optical package, only relevant to the cable space. While some deviation has come from the typical butterfly "analog" laser package in the last few years, the gains have been minimal. Also, legacy DMTx have the necessity for electronic harmonic distortion correction, both for residual CSO and CTB from the analog laser, and for fiber induced CSO, ultimately the extent to which these corrections are utilized also creates a power consumption and size threshold in the electronics.

In the case of externally modulated transmitters (EMTx), where they are highly desired for their low noise capability and lack of high CSO accumulation over fiber, their size, power draw and price has typically made them unattractive in comparison to DMTx. For an EMTx, a high power CW laser, external modulator and SBS suppression circuitry typically define the size and power consumption.

The ability to simplify the technical requirements for CRM allows for the concept of creating a small form factor pluggable transmitter. There are two key requirements: the first is finding the smallest possible optical component packages that would be able to meet CRM linearity requirements and the second is to collapse the new necessary RF electronics into integrated circuitry. These two steps create a framework under which one could envision smaller packaged transmitters.

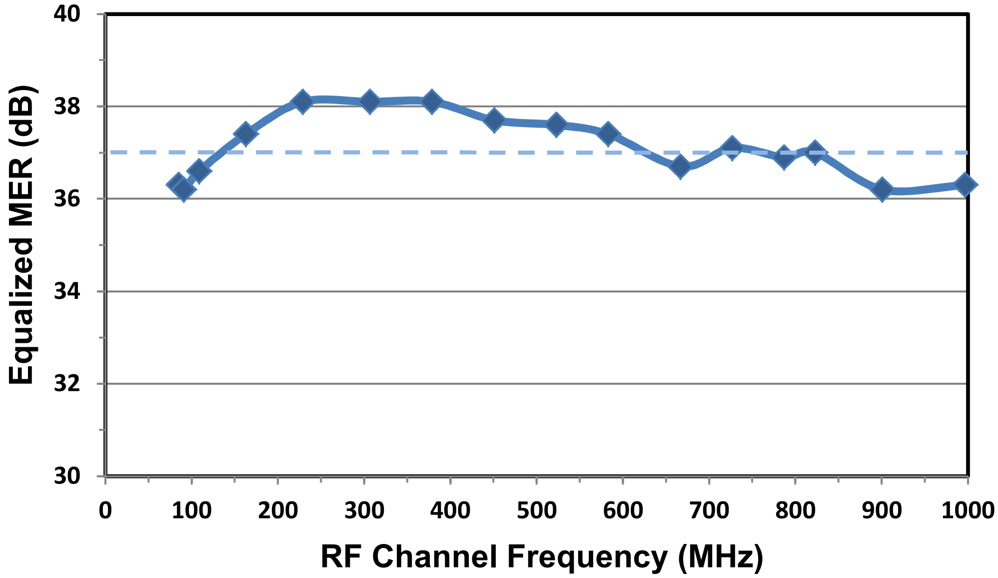

Figure 2 – Experimental Results for a Rudimentary, Proof-of-Concept XFP Transmitter

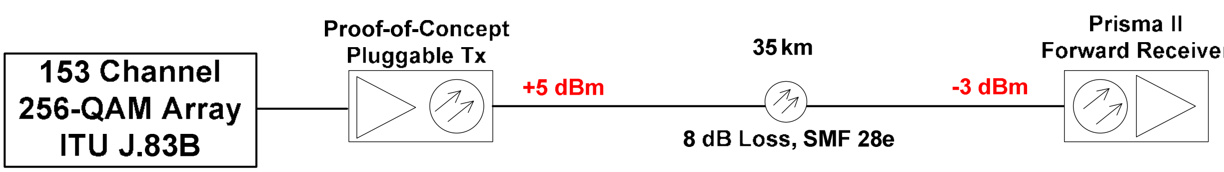

Figure 3 − Optical Test Link

<u>Proof of Concept</u>

Figure 2 shows the equalized MER performance of a rudimentary proof of concept transmitter whose optical package and pertinent RF electronics can fit comfortably in a 10 Gigabit small form factor pluggable XFP package. Figure 3 depicts a representation of the test link set up, where the proof-of-concept transmitter with an output of 5 dBm is followed by 35 km of fiber into a Prisma II forward receiver with an input power of -3 dBm. The input channel loading was 153, 256-QAM, ITU-T J.83 Annex B channels. The channel loading spanned 1 GHz to 82 MHz, leaving room for expected growth in the return path. We expect that integration and optimization of components will improve both the performance and reach.

<u>A Standards-Based Pluggable Transmitter: An XFP Form Factor CRM Transmitter for the Cable Space</u>

A question arises as to what would constitute an acceptable small form factor package, keeping in mind that the cable industry wishes to adopt, where possible, industry standards for form, fit and function without hampering the competitive aspects of functionality and added value. In the baseband digital space there are various multi-source agreements (MSAs) for small form factor transceiver packages. For example, XENPAK, 10 Gbps small form factor pluggable (XFP), small form factor pluggable (SFP), and small form factor pluggable 10 Gbps (SFP^+) are all popular such multi-source agreements.

MSAs are by definition industry-accepted by producers and users; they have been in place for some time, benefitting from mature components and predictable cost reduction curves. Note that the majority of the MSA documentation is focused on interface standards such as powering, signaling, monitoring, DC power limits and thermal dissipation, physical outline and mechanical specifications. It is absolutely conceivable that the cable space can leverage such an MSA interface, while internally maintaining functionality specific to CRM operation.

As an example of such a usable standard see Figure 4 which details the physical connector interface for a 10 Gbps XFP, as defined in the XFP MSA, [8.] The XFP specification, as currently defined, gives a CRM transmitter the opportunity to use several already existing, industry-standard interface lines (those NOT highlighted in a red circle) for powering, communication, control, and modulation inputs. For example, the differential input signal interface specified for the XFP (TD^+ and TF- on pins 29 and 28) lend themselves well to RF QAM input. Various powering options exist, including 1.8 V (Vcc2), 3.3 V (Vcc3), and 5 V (Vcc5), with nine ground pins specified for excellent RF and DC connectivity.

Module control is equally straightforward with an industry-standard two-wire serial interface (serial clock, SCL, and serial data, SDA), along with interrupt, module de-select, power-down, module numbering and presence detection. Communication protocols to the pluggable module are called out in the XFP specification, along with allowable DC dissipation limits for each power line and for the module as whole. One difference between the XFP as specified in the MSA and the transmitter under discussion here is that we refer to a transmitter, not a transceiver. Only the forward path transmitter will be present; the return path can be located separately. This releases 5 pins specific to the receive portion of a transceiver, received data (RD+, RD-), reference clock (REFCLK+, REFCLK-), and receiver loss of signal (RX_LOS), further easing host PCB layout.

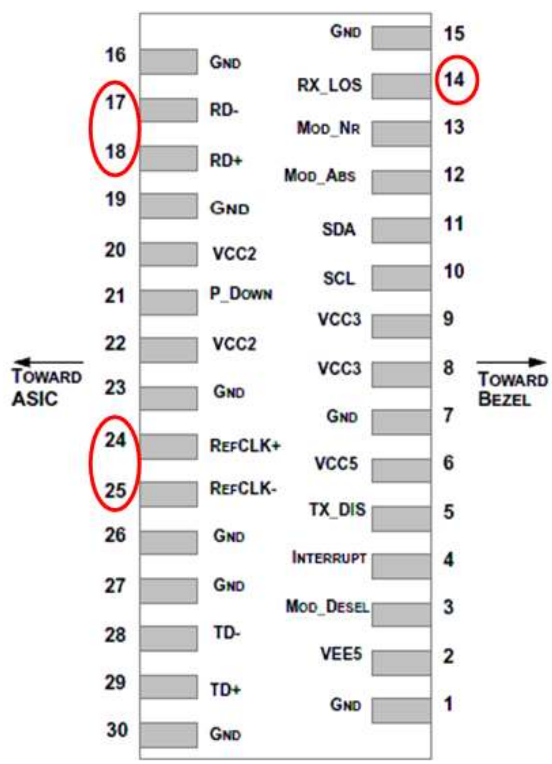

Figure 4 –XFP Interface: MSA-Defined Specification

Transmitter Functionality Options

Regarding the operation of small-form-factor transmitters, there are two options that have been expressed so far. The first is where the functionality of the pluggable transmitter is totally self-contained. This means that the burden of linearity for the transmitter is internal and independent of the incoming signal. The only interaction with its host would be for powering and status monitoring. As is the case now with most, if not all, contemporary HFC platforms, where an RF broadband signal comes in to a transmitter, it is analyzed and corrected for distortions and optically modulated to exit. The second option is that the transmitter linearity is not self contained, and is dependent on an input signal that has already been analyzed and corrected for distortions, leaving the pluggable transmitter to have only the function of optical modulation [9].

While both options are technically feasible and have their advantages and disadvantages, we prefer the self-contained transmitter approach for various reasons: It allows for a potential standardization based purely on interfaces and not tying performance parameters across active components, beyond what exists already at the RF level, such as the DOCSIS Downstream Radio Frequency Interface Specification (DRFI). It frees potential higher layer host platforms from having to carry calibration data for pluggable optics. And finally it is fundamentally important that host platforms not be transmitter-specific, that is, DMTx, EMTx, 1310 nm, DWDM, etc. Conversely, the option that it is not self-contained could possible lead to a better cost structure up front, though not including deployment issues involving cross vendor reliability and traceability of faults when failures occur.

IT'S EASY BEING GREEN

The XFP MSA specification [8,] defines the maximum DC power dissipation levels for the XFP package. The DC dissipation for a Power Level 3 device is specified at not more than 3.5 W per module. *We propose adopting this dissipation limit for RF optic pluggable transmitters*. At a 3.5 W maximum DC dissipation per optical transmitter, the DC power dissipation reduction from current, state-of-the-art transmitters which consume anywhere from 7 – 15 W, is a very green ***50 to 75% per transmitter***.

As large as the power consumption reduction is, it still does not completely capture the power savings due to status monitoring and control being directly absorbed into the content generation hardware, CMTS, or RF gateway. No external processors or aggregators are necessary to control the pluggable XFP, outside the content generation hardware, whether it be a CMAP box, a CESAR router, or a CMTS. Optics are now truly part of the "smart box" with high-speed backbone inputs and access optical outputs.

Existing headend architectures still create channel lineups *manually* via lossy RF combiner structures which aggregate content generation outputs before applying them to the access optics. As content generation evolves, generation devices which synthesize entire channel lineups via direct digital synthesis (DDS) are becoming available. Such devices can or will generate 135 to 153 channels per port for direct transmission to the node. Such "full-spectrum" port synthesis creates an additional opportunity for a significant power savings within the content generation boxes themselves, *on the order of 3-5 watts per port*, by reducing the port's output levels generated when attached to a pluggable optical module.

Programmable port power can reduce some >350 W per 96-downstream signals per port device, just by leveraging the lower output levels required of an RF XFP directly driven by a content generator's output port. Under the tightly integrated control of the content generator, the XFP also becomes part of an "agile" channel lineup, allowing sparing, redundancy, and idling of unused functionality to optimize power consumption versus bandwidth requirements. Whether by itself or integrated into the content generator, an XFP-based RF optical transmitter can save between 5 to 12^+ watts per port in DC power dissipation. Considering an even optimistic 85% off-the-wall efficiency improvement yields a 6 to 15^+ watts PER PORT savings.

CONCLUSIONS

We have shown that pluggable RF optics within a next generation access IP platform are possible, without triggering changes to the outside plant infrastructure. Use of an all-QAM complex RF modulation payload allows simplified, relaxed design rules in the access link to the customer. We use this to create a new means to recapture space, power, and cost by the use of specification-based pluggable optical transmitters. Pluggable RF XFP optical transmitters move the electrical to optical transition from a separate chassis to an integral part of the IP platform, tightening network control, lowering total power requirements, and saving rack space.

Acknowledgement

We would like to extend our gratitude to John Chapman, Ron Hranac, Steve Condra, Hisham Saad and Dan Ott of Cisco, for their support and helpful discussions.

References

[1] Chris Bowick, *"The Winning Network,"* Multichannel News, November 1, 2010.

[2] Steve Condra, Discussion, Director of Network Architecture, Cisco.

[3] Arris Group Inc., *"IP Video…An Architecture in Transition,"* Presented to the Greater Chicago SCTE Chapter, May 2010, http://www.gcscte.org/presentations/2010/IP%20Video%20arch%20Chicago%20SCTE.pdf

[4] Sheryl L. Woodward, Mary R. Phillips, *"Optimizing Subcarrier-Multiplexed WDM Transmission Links,"* Journal of Lightwave Technology, Vol. 22, Nr. 3, March 2004.

[5] Gèrard Terrault, *"QAM Signal Impairments And Their Effects on MER and BER,"* Sunrise Telecom, White Paper, Rev 1.04, 2-Jan-03, http://docsis.beckitrue.com/documents/sunrise/QAM_Impairment Effects_on_MER_BER_104.pdf

[6] Ron Hranac (Cisco) and Bruce Currivan (Broadcom), *"Digital Transmission: Carrier-to-Noise Ratio, Signal-to-Noise Ratio and Modulation Error Ratio,"* Cisco/Broadcom, White Paper, http://www.cisco.com/en/US/prod/collateral/video/ps8806/ps5684/ps2209/prod_white_paper0900aecd805738f5.html

[7] T. Brophy, F. Villarruel, K-Y. Wu, *"The Delicate Balance,"* Cisco, White Paper, http://www.scientificatlanta.com/products/customers/white-papers/g1667a_thedelicatebalance.pdf

[8] XFP, MSA, ftp://ftp.seagate.com/sff/INF-8077.PDF

[9] David Piehler, *"Next-Generation Components for Optical Access Networks,"* Neo-Photonics USA, Optical Fiber Conference Proceedings, 2011.

Author's Contact Information
Fernando Villarruel: villarf@cisco.com
Fred Hirt: frhirt@cisco.com
Kuang-Yi Wu: kuanwu@cisco.com