PennWell®

LIGHTWAVE®

EDITORIAL GUIDE

Emerging 100G applications and technologies

It's clear that the deployment of 100-Gbps technology is well underway. But that doesn't mean that developments have stopped. These articles illustrate several paths toward next-generation 100G systems and networks.

SPONSORED BY:

Alcatel·Lucent

Originally published May 1, 2013

100G reaches the metro

by **STEPHEN HARDY**

METRO NETWORKS WEREN'T the first applications of 100-Gbps technology. But operators within the niche appear to be making up for lost time.

As 100-Gbps deployments in long-haul networks gained momentum last year, several observers thought a less expensive alternative would have to be offered for 100 Gbps to see widespread deployment in the metro, particularly for data-center interconnect.

Events this year have proven such fears unfounded. This year's tide of 100-Gbps deployments includes several examples within metro/regional networks. But that doesn't mean those who called for a different technical approach to the metro were wrong.

Pent up demand

The increasing number of metro/regional 100-Gbps deployments highlights how quickly bandwidth demand has grown across network niches.

"Metro 100 gig has ramped quite well for us. About half of our customers that are deploying 100 gig are deploying 100 gig in the metro," says Michael Adams, vice president, product and technology marketing at Ciena Corp. The company counts CenturyLink and XO Communications among its high-profile metro 100G customers.

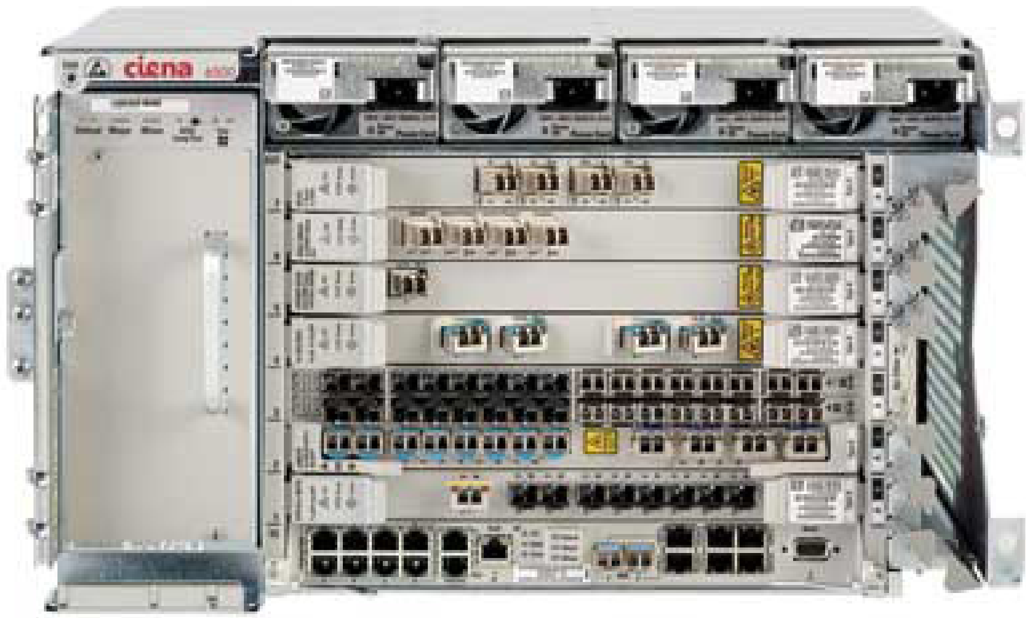

Systems that use coherent detection have proven almost as popular in metro applications as they have in the long haul. Photo courtesy of Ciena Corp.

Ciena has found a healthy demand for both direct 100 Gigabit Ethernet (100GbE) connections - mainly for router interconnect, particularly from the larger carriers - and 10×10-Gbps muxponding. “Right now I would say almost about 40% of our deployment is with 100GbE, in the context of running them over 100-gig waves,” Adams explains. “Now that’s skewed a little bit to a smaller number of customers. I would say from a number-of-customers perspective, the statistic would be more skewed to 10×10G.”

Mike Sabelhaus, who works on optical solutions for Fujitsu Network Communications, reports that 100G technology has appealed to several different carrier types. For example, one “large east coast wireless network operator” is using the technology to connect routers for mobile backhaul, he says. Wireless operators in general have adopted 100G for connecting mobile switching centers, Sabelhaus adds. Elsewhere, another large carrier customer has incorporated 100G technology into its network to meet a trio of requirements: fiber relief via 10G muxponding, IP router aggregation/interconnection, and support of Optical Transport Network (OTN) switching requirements.

Cable operators have interest in the technology as well. “A lot of what we’re seeing in the cable networks is using it for their metro transport networks, hub site to hub site within a city or region, and potentially to connect backbones together or other regions together," Sabelhaus notes.

Like his competitors at Ciena, Sabelhaus sees significant interest in both direct 100GbE connections and muxponding applications. That trend should continue in the future, he says, based on his projections for demand between this year and 2015. "In the transponder area, I'm seeing roughly 50% to 75% growth from year to year," Sabelhaus reveals. "And muxponder is in the 30% to 40% growth area."

Demand for high-speed metro/regional pipes drives not only sales of 100G platforms, but also 100G network services. Zayo Group has just started to deploy its first customer-requested 100-Gbps metro link after building a metro-centered extension of its long-haul network for another customer.

The drivers for 100-Gbps network services are similar to those the systems vendors have encountered, reports Zach Nebergall, head of the wavelength product unit at Zayo. "Data center to data center is I'd say the primary driver we've seen so far," Nebergall says. "In the carrier segment, I'd say there is a similar dynamic, but a PoP [point of presence] to PoP scenario; with the wireless guys, it's a mobile switching center to PoP."

Coherence around coherent

Another commonality among the majority of metro/regional deployments is use of coherent transmission technology. Part of this trend may be that only one optical hardware supplier - ADVA Optical Networking - has commercialized an alternative that uses direct detection (more on this later). Nevertheless, coherent has proven more attractive than the doubters had believed.

Helen Xenos, Ciena's director of product and technology marketing, asserts that from a technical standpoint, coherent is popular for metro/regional applications for the same reasons it's popular in long-haul networks. It works on just about any fiber, is spectrally efficient, and delivers robust performance in the face of impairments.

Comfort with the technology also has proven a major factor with some deployers. "Given that it's early in the deployments, we think that the cost basis was still there, that we could make the same technology work," Nebergall says of Zayo's decision to use coherent technology. "The tradeoffs didn't make sense to us such that we'd move off of a piece of equipment that we were pleased with so far."

Yet, just because you're using the same system in your metro/regional and long-haul networks doesn't mean you're using exactly the same technology. For example, Adams acknowledges that Ciena has a portfolio of coherent combinations for different applications. While declining to provide a great deal of detail, he reveals that Ciena supports more than one approach to forward error correction so customers can balance reach and latency requirements. "We've talked about in the past that for metro, as an example, one possibility going forward is using 16QAM tuning capability to have a 50% economic improvement in metro/regional 100 gig," Adams adds.

Direct detect for the metro

But coherent hasn't succeeded within every corner of the metro. ADVA Optical Networking has found several customers for its 4×28-Gbps direct detect approach. In particular, cloud service providers, Internet exchanges (such as the Amsterdam Internet Exchange and DE-CIX), and private enterprises find direct detect appealing, according to Jim Theodoras, senior director of technical marketing at the systems house.

While the technique offers a smaller price tag than coherent, Theodoras says that factor often appears to be beside the point. "We were surprised at how less important 'cost competitive' was in this space," he says. "It turned out they cared more about compatibility and efficiency."

ADVA Optical Networking reports success with its 4×28-Gbps approach among those outside of the traditional telco establishment. Photo courtesy of ADVA Optical Networking

Among these customers, "compatibility" means dovetailing with existing 10-Gbps traffic and networks; it's comparatively easy to move from 10-Gbps to 25-Gbps wavelengths, particularly on the same network, Theodoras says. And "efficiency" means small power and footprint requirements as well as greater spectral efficiency versus 10 wavelengths of 10 Gbps.

Of course, four 25-Gbps wavelengths don't provide the spectral efficiency of a single 100-Gbps wavelength. Yet the gap between 4×45G and 100G may not be as great in some applications as it would appear, Theodoras asserts.

"It turns out that a lot of the existing 100-gig coherent offerings, to meet the type of span requirements [some users require on brownfield deployments], they need guard bands on each side of it," he explains. "So if you look at the guard bands required in a 100-gig coherent versus the 100-gig metro, it turns out it's not really a 4:1 disadvantage for [4×28-Gbps] 100-gig metro versus the coherent; it's almost a 2:1 disadvantage if you take into account the guard bands."

Cost does become an issue on the other side of the span spectrum, where applications run relatively short distances. Both of the Internet exchanges, for example, needed to send 100 Gbps less than 40 km. "If someone only needs to go 100 km, they don't want to pay for a 3,000-km solution," Theodoras reasons.

As this article is being written, ADVA Optical Networking is the only systems house to offer the 4×28-Gbps option as a publicly released product - but others, such as ECI Telecom, have signaled their intentions to join the party.

The question, given the success coherent approaches have had in the metro, is just how large that party will be.

STEPHEN HARDY is Editorial Director & Associate Publisher of *Lightwave*.

Optical Networks Light Up Broadband Experience

By David Stokes

I was fascinated to hear Jim Cramer of TV's Mad Money say: "One word: Optical," on a **recent episode**. Cramer was referring to the transport networks required to support high performance video. His observa[illegible] acknowledges the reality known by te[illegible] operators for some time now; consum[illegible] are using new devices that offer new services that fuel an exponential need [illegible] more bandwidth. Telecom operators [illegible] to transform their network to cope with these new demands in a cost effective way. The optical network is at the heart of supporting this transformation by providing high capacity connectivity along with the network flexibility and intelligence required to maximize network utilization.

Telecom operators are embracing the need for a scalable, versatile and dynamic networking solution that can adapt to the ever-changing nee[illegible] users. That's evident with cus[illegible]ers [illegible] **Turk[illegible]** **Superonline**, **China Uni[illegible]** [illegible] **[illegible]OR[illegible]** They have all made th[illegible] that give them the capacity to support ultra-broadband, as well as the ability to reduce operational costs and increase revenue.

> Read about **Record-Breaking [illegible] in a 100G Trial with Türk Tele[illegible]**

NETWORK CAPACITY: A MAJOR CHALLE[illegible]

Consumers are driving [illegible] more bandwidth; they [illegible] using bandwidth hung[illegible] and services supported on a new era of connected devices such as smartphones, smart TVs and tablets. We are [illegible] telecom operators in a race t[illegible] their networks up to speed [illegible] consumer demand. Those [illegible] fall [illegible] risk losing their high value clients. New services like IP Video are driving capacity growth in fixed networks, where we see traditional broadcast TV companies, satellite TV delivery companies and Pay TV companies all introducing IPTV services to complement their standard delivery mechanisms. In the mobile space we see LTE rollout rapidly driving up mobile data usage, and the backhaul network must [illegible]dwidth [illegible] ent[illegible]pris[illegible] [illegible]ased se[illegible] dri[illegible]ment [illegible]r high-[illegible] late[illegible]cure connecti[illegible]

We see [illegible] slow-down in this growth with research suggesting that global IP traffic will increase threefold over the next three years and that mobile data traffic will increase 13-fold between 2012 and 2017. This growth is being driven by the widespread rollout of LTE, rapid expansion of cloud services and new higher capacity TV formats such as 3DTV and 4K UltraHD.

> Read why high-volume **metro market will drive 100G to mass market volumes**

[illegible]TM[illegible]T IN TH[illegible] [illegible]SP[illegible]T NETW[illegible]K: [illegible] MOV[illegible]IGHT

Telecom operators are seeing the need for a new investment cycle in the transport [illegible]work to [illegible]ort this [illegible] uplift in [illegible] evolve this to 400G in the future.

> **Listen to the podcast** to learn more.

> Find out more: Listen to our **podcast** and read the **Turkcell Superonline**, **China Unicom** and **NORDUnet** case studies.

> To contact the author or request additional information, please send an email to **techzine.editor@alcatel-lucent.com**.

We also see that telecom operators are looking for far more than just bigger pipes. They want a scalable, versatile and dynamic networking solution that can adapt to the ever-changing needs of their users, drive do[illegible] cost of ownership and allow oper[illegible] re[illegible]lize the full potential of their netw[illegible] investment. The network must pro[illegible] flexibility to support driving down costs, the intelligence to easily add new services and the agility to scale bandwidth [illegible]ere and when it is needed. At Alcatel-[illegible]ent we call this **Agile Optical Networking** [illegible]ure 1).

Figu[illegible] Alcatel-Lucent [illegible]e Optical Netw[illegible]ng summary

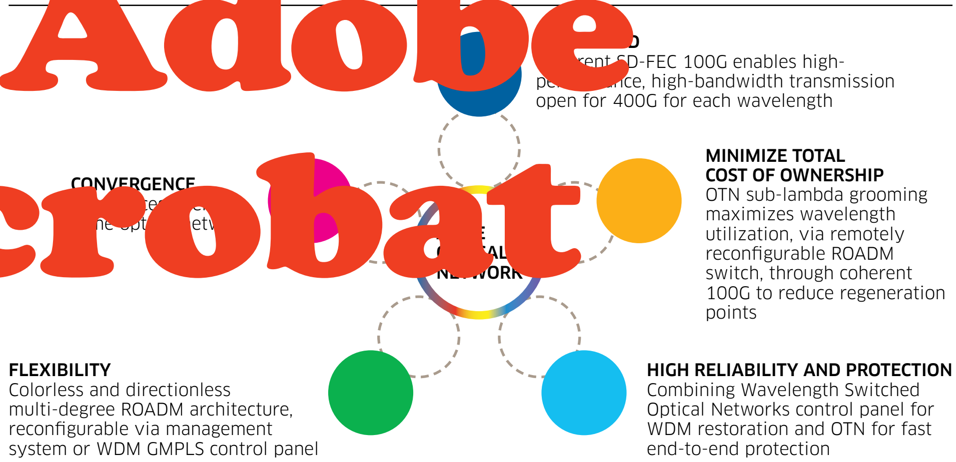

Originally published November 4, 2013

Time-domain pulse shaping for increased spectral efficiency

by **STEPHANIE MICHEL,** Agilent Technologies

IN OUR PREVIOUS articles (see "Why Complex Modulated Optical Signals?," "Complex Coding Concepts for Increased Optical Bit Transfer Efficiency," and "Which Optical Modulation Scheme Best Fits My Application?"), we got to know coding and modulation techniques to maximize the number of bits transferred per unit time over a fiber. We saw that using the phase of an optical signal for encoding information, in addition to other parameters, can multiply the data capacity of a channel.

In the pursuit of using existing limited bandwidth resources most efficiently, we have to be aware that the complex modulated signal spread in time and that consecutive symbols may overlap, which is referred to as inter-symbol interference (ISI). ISI induces errors in the interpretation of the signal on the receiver side. In the frequency domain as well, we have to take care to avoid interference between adjacent channels. For data rates of 400 Gbps and 1 Tbps, this will become a topic of crucial importance.

This article elaborates on the conditions needed to prevent these effects and on the use of different filtering techniques for the purpose of bandwidth and signal containment.

Nyquist ISI criterion

The Swedish engineer Harry Nyquist explained in the 1920s that for eliminating ISI, the impulse response *h(t)* needs to fulfill the following requirement in the time domain:

$$h(nT_S) = \begin{cases} 1, & n = 0 \\ 0, & n \neq 0 \end{cases}$$

for all integers n. T_S is the pulse spacing of adjacent pulses.

In Figure 1, a signal that meets this condition – sinc *(t)* pulses – illustrates the impact of the criterion: The pulses overlap, but only the sampled symbol contributes to the response at the sampling instant t_k. The other symbols are zero at this point in time. In this way, we have avoided the signal-degrading and error-inducing effects of ISI.

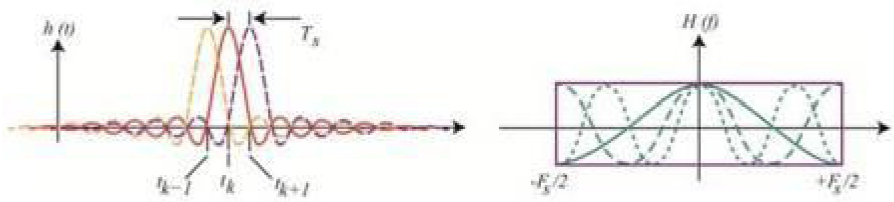

FIGURE 1. *Orthogonal sinc (t) = (sin(t))/t pulses meet the Nyquist ISI Criterion. [2]*

The right side of Figure 1 shows the Fourier transform (FT) of the impulse response. You can see that the frequency response that fits into a rectangular frequency window fulfills the Nyquist ISI criterion:

$$F_S \sum_{k=-\infty}^{+\infty} H\,(f - k \cdot F_S) = 1 \qquad \text{with} \qquad F_S = \frac{1}{T_S}$$

This means that the harmonics–the components with a frequency that is an integer multiple of the so called Nyquist frequency F_S–must add up to a constant value, in order to fit into a frequency band without ISI. The Nyquist frequency F_S is the minimum bandwidth needed to encode a signal without loss of information.

Nyquist pulse shaping using a finite impulse response filter

A sinc-signal may be perfect for preventing ISI as described above, but it is not practical because it is infinitely extended in time. Therefore we need to truncate it in the time domain, which is performed in practice by using a finite impulse response (FIR) filter. An FIR filter of order *R* responds for *R*+1 sample points and then returns to zero. Only past sample points *x[n-i]* are considered for the convolution of the filter output *y[n]* so that filtering can be performed in real-time.

The discrete-time FIR filter output *y[n]* in dependence of the input *x[n]* is described by:

$$y[n] = \sum_{i=1}^{R} b_i \, x[n-i]$$

where b_i are the filter coefficients.

To avoid aliasing, the pulse-shaping FIR filter must oversample by a factor of q=2 at least. In other words, there must be at least 2 sample points within T_S, so that the pulse shape can be reconstructed on the receiver side, without losing higher-frequency components.

Figure 2 shows the filter results for a sinc pulse filtered at different filter order *R*, always oversampled by a factor of q=2. The power spectrum results from the convolution of the rectangular shaped spectrum of the sinc-pulse with the sinc-shaped spectrum of the rectangular window.

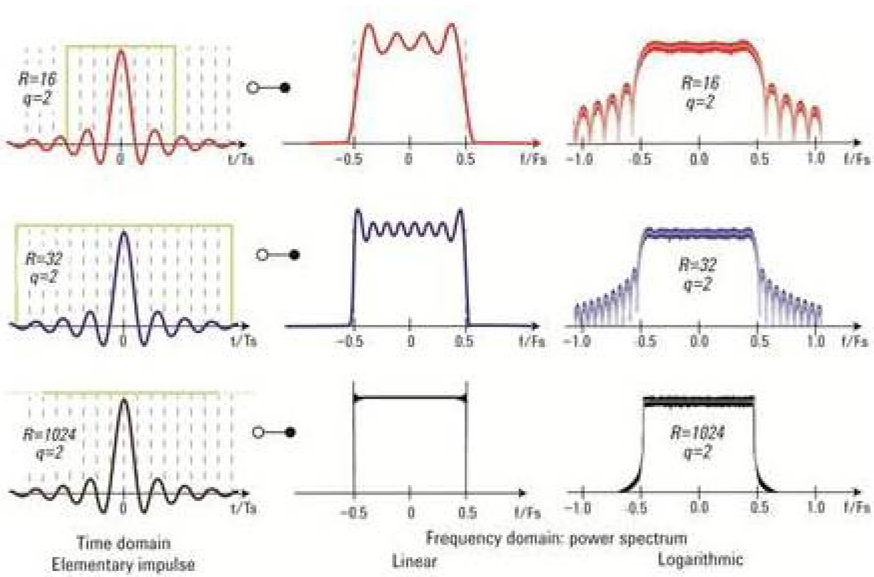

FIGURE 2. *FIR filter of different order R used to truncate sinc-signal: time domain waveform, after Fast Fourier transformation power spectrum on linear and on logarithmic scale. [2]*

In the first line, with a filter order of 16, the signal spans 8 T_S. In the FFT, we can see distortions caused by the limited time window. Most of the power lies within the Nyquist band (-0.5 F_S to 0.5 F_S) but a part lies outside the band. The power spectrum shows the harmonics as image spectra.

If we double the length of the filter (*R*=32), the signal better fits into the bandwidth, but we still see some ringing. At *R*=1024, the spectrum looks almost ideal; the ringing is only visible at the steep edges and the power spectrum also shows less out-of-band contribution. Unfortunately, the higher the order R of the filter, the higher the complexity of the filter design too. Therefore, you would usually stay with the lowest R that answers the requirements.

Raised cosine filters

For even better out-of-band suppression and a spectrum without ringing, raised cosine filters provide a reasonable alternative. Here, the impulse response is dependent on the so-called roll-off factor α which can take any value from 0 to 1:

$$h(t) = sinc\left(\frac{t}{T}\right)\frac{\cos\frac{\pi\alpha t}{T}}{1-\frac{4\alpha^2 t^2}{T^2}}$$

The raised cosine filters also fulfill the Nyquist ISI criterion that only the sampled symbol contributes to the signal. All other symbols are zero at the sampling points. In comparison to sinc-shaped pulses, the raised cosine signals require more bandwidth.

Figure 3 depicts the filter response for four different roll-off factors α:

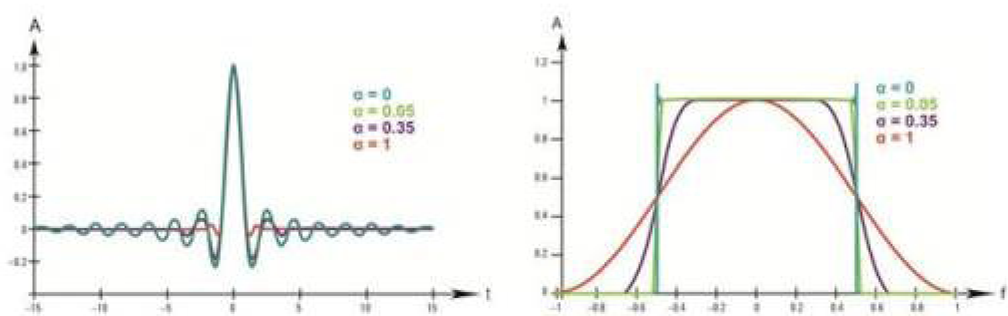

FIGURE 3. *Raised Cosine Filters with different roll-off factors α: normalized time and frequency domain presentation. [2]*

In the frequency response, it stands out that for any value of α, the curves are crossing the same point at $\pm F_S/2$, which is half the pulse rate. As mentioned before, this is the Nyquist frequency, the minimum bandwidth needed for a data transfer without loss of information. Apart from that, you can see that for α = 1, we hardly have any ringing but the frequency spectrum does not fit into the bandwidth. For α = 0, it is vice versa: the frequency response is rectangular within the bandwidth (the overshoots at the edges are a mere mathematical effect, also known as the Gibbs Phenomenon, and do not have any practical impact). The time domain signal now shows more ringing, though.

Why should we care about the ringing when we first stated that at the sampling point only the sampled symbol contributes to the signal? In practice, ringing is a problem because the other symbols have zero contribution when sampling only at this ideal instant. Under real world conditions, it is almost impossible to have the receiver sample exactly at this point in time so that we'll always have some ISI that induces errors in the interpretation of the signal.

So, we obviously have a tradeoff between bandwidth containment and suppression of ringing in the time domain. A compromise in the choice of an adequate value of α has to be found on an individual basis for every fiber-optical application.

Raised cosine filters in practice

Let's have a closer look at the influence of different roll-off factors on the most promising modulation scheme for 400 Gbps and higher data rates: 16-QAM. Figure 4 shows measurements not only of the frequency domain response but also the eye diagrams and the impact on the transitions between constellation points and the constellation points themselves.

The top example shows the case of non-shaped rectangular pulses. As we know, a signal that only occupies a fixed time interval has an infinitely extended frequency spectrum; we can see in the frequency response large side slopes. The eye diagram shows the typical behavior of a wide-band signal with open eyes. Between the constellation points, we have straight transitions.

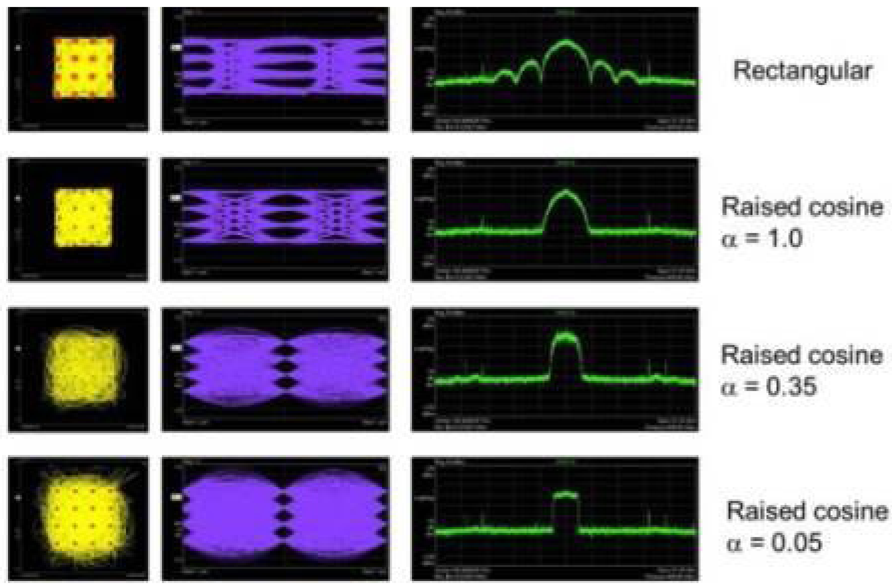

FIGURE 4. *Raised cosine filters on 16-QAM signal in dependence of roll-off factor α: constellation diagram, eye diagram and frequency spectrum; signals created with an Agilent M8190A Arbitrary Waveform Generator.*

Using a raised cosine filter with a roll-off factor $\alpha = 1$, the frequency spectrum becomes narrower; the side slopes are not visible anymore. The eye diagram shows wide-open eyes. The constellation points are smaller. This is typical for a system with reduced bandwidth. The detection bandwidth on the receiver side is, by implication, also reduced, which lowers noise.

At a roll-off factor $\alpha = 0.35$, the frequency width has further decreased, and with it the size of the constellation points. The transitions between the constellation points start to show much overshoot. This is because when reducing bandwidth the transitions between the symbols get extended in time, which is reflected in the constellation diagram by the long curved transitions between the points. The eyes are closing and therefore the sampling timing gets more critical.

An almost perfect rectangular spectrum is reached at $\alpha = 0.05$. The transitions between the little constellation points show large overshoot. The completely closed eyes indicate that for avoiding errors, the sampling point has to be adjusted precisely.

How much spectral efficiency do we gain?

To get an idea of the quantitative gain in spectral efficiency by pulse-shaping filters, let's compare it to the effects reached by applying orthogonal frequency-division multiplexing (OFDM).

Figure 5 gives a brief recap of the OFDM principles, which are similar to the Nyquist format.

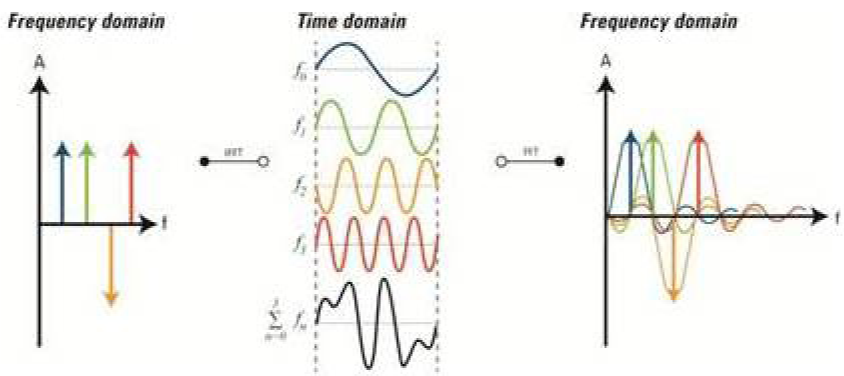

FIGURE 5. *OFDM in the frequency and in the time domain. [2]*

In OFDM, the frequency sub-spectra are sinc-shaped. For increased spectral efficiency, they are overlapping but because of their orthogonality – meaning that they are shifted by multiples of π/2 – they do not interfere with one another. In the time domain, a symbol is a sum of sine curves with equidistant carrier frequencies fn in a fixed time window. In this example, we have in one channel four subcarriers at four frequencies. The orange is phase shifted by π, as can be seen after inverse Fast Fourier transformation (IFFT).

Figure 6 now shows spectral analysis on a measured 16-QAM modulated OFDM signal.

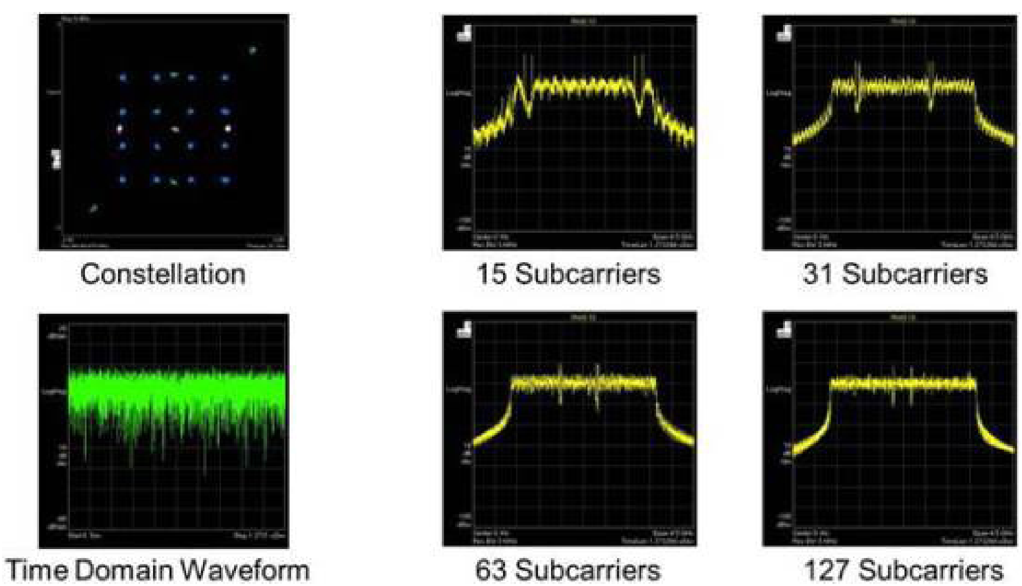

FIGURE 6. *Impact of OFDM on a 16-QAM signal spectrum in dependence on the number of subcarriers; signals created with an Agilent M8190A Arbitrary Waveform Generator.*

On the top left side, you can see the constellation diagram and below the time domain waveform. With 15 subcarriers and two pilots, we get a fairly flat frequency spectrum and a steep roll-off.

By increasing the number of subcarriers, the spectrum is flattening and the two pilots are moving towards the center. At the bottom line, we see that the spectrum is approaching rectangular shape with the growing number of subcarriers.

How does this translate now into a gain of spectral efficiency in comparison to that reached by Nyquist pulse shaping? For the sake of comparability, in Figure 7, the normalized spectral efficiency (SE) is plotted over the Nyquist filter length *R* (the oversampling factor *q* is chosen again to be 2) versus the number of OFDM subcarriers *N*.

The figure shows that both techniques provide almost the same result regarding spectral efficiency.

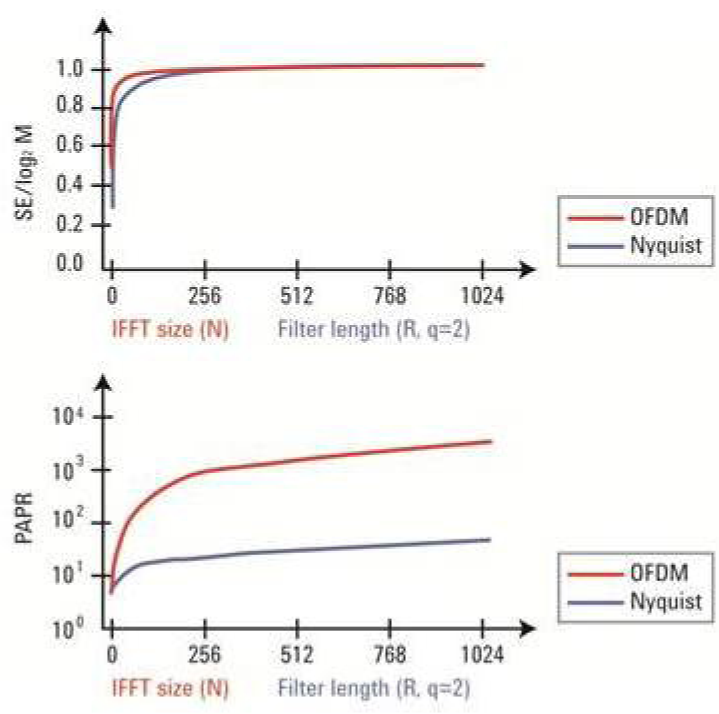

FIGURE 7. *Effect of Nyquist pulse shaping versus OFDM on spectral efficiency and peak- to-average power ratio (PAPR) [2].*

The comparison of the also normalized peak-to-average power ratios (PAPR) reveals similar behavior at differing levels. The PAPR of OFDM time domain waveforms is much larger. This phenomenon owes to the fact that in In OFDM, the signal exhibits some peak values high above the average power level. As a consequence, OFDM circuits and test instruments need a large dynamic range to avoid distortions induced by clipping the higher levels. For optimized spectral efficiency, pulse shaping can be applied in combination with OFDM.

Up to this point, we have touched on the technical requirements for the more sophisticated data transmission concepts several times. In the next article of this series, we'll investigate in detail the challenges imposed by complex modulation for the technical implementation of an optical transmitter.

Figures: Dr. Bernd Nebendahl, René Schmogrow

References

1. H. Nyquist: Certain Topics in Telegraph Transmission Theory, Trans. AIEE, Vol. 47, pp. 617-644, Apr. 1928
2. R. Schmogrow M. Winter, M. Meyer, D. Hillerkuss, S. Wolf, B. Baeuerle, A. Ludwig, B. Nebendahl, S. Ben-Ezra, J. Meyer, M. Dreschmann, M. Huebner, J. Becker, C. Koos, W. Freude, and J. Leuthold,: Real-Time Nyquist Pulse Generation Beyond 100 Gbit/s and its Relation to OFDM, Optics Express, Vol. 20 (1), pp. 317 – 337, Jan. 2012

STEPHANIE MICHEL is technical marketing engineer in the Digital Photonic Test Division of the Electronic Measurements Group at Agilent Technologies.

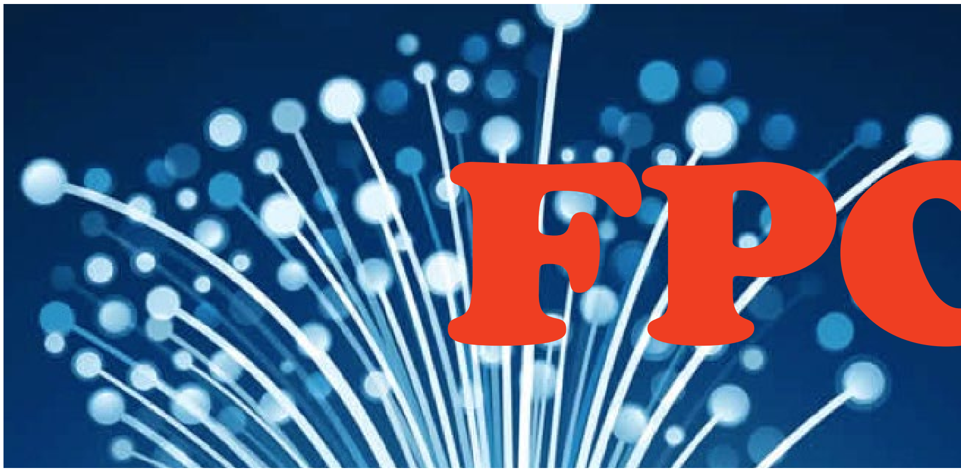

- Read the Press Release: **Alcatel-Lucent and BT achieve fastest real-world fiber speeds of 1.4Tb/s with a world record spectral [illegible]ficiency of 5.7b/s/Hz over core network**
- [illegible]ow **BT Ingenious News**
- Learn more about the Alcatel-Lucent **Agile Optical Networking** solution.

BT and Alcatel-Lucent achieve fastest ever real-world fibre speeds

BT and **Alcatel-Lucent** have be[illegible]g on a ground-breaking field tria[illegible]t has recorded fibre speeds previously unconquered in the real-world.

Using the latest technology, fibre speeds of 1.4 Tbits/s were achieved on real fibre cable outside of the research lab environment. It is believed these spee[illegible] are the equivalent of transmitting 44 uncompressed HD films in a single second.

A pioneering 'flexible grid' inf[illegible]uc[illegible] (Flexgrid) was used during the [illegible] t[illegible]y the gaps between transmissio[illegible] Flexgrid, backed by **Bell Labs technology innovations**, maximises the efficiency of the existing installed fibre, meaning t[illegible] is no need to put in costly new fibre[illegible]

PUSHING BOUNDARIES

Results showed that by reducing the spectral spacing between the channels from 50GHz to 35GHz, spectral efficiency is enhanced by almost 43 per cent compared with today's network standards. A record spectral efficiency of 5.7b/s/Hz was achieved – which is equivalent to fitting a 1T Super-Channel in less than 200GHz spectrum.

This shows Flexgrid is key to creating high-capacity, spectrally efficient Super Channels.

[illegible] he[illegible] of researc[illegible] said: "[illegible]story of [illegible]ng innovation in telecommunications, from the earliest days of the electric telegraph to today's global fibre networks. These trials co[illegible]ue that [illegible]tion as [illegible] work w[illegible] Alc[illegible]nt [illegible] the ever increasing bandwidth required by our customers, and deliver new and exciting serv[illegible] which rely [illegible]st, data-hungry [illegible]ation[illegible]

ALIEN [illegible]

These revolutionary speeds were recorded by trialling an existing 410km fibre link between the BT Tower in Londo[illegible] BT's A[illegible]

Th[illegible]duct[illegible] an "Alien Super Channel" made of seven 200 Gbit/s channels bundled together to provide a combined capacity of 1.4 Tbits/s.

The Super Channel is "alien" because it operates transparently on top of BT's existing optical network.

MEETING THE NEEDS OF AN EVOLVING MARKETPLACE

The trial demonstrated that using the Flexgrid approach can increase BT's existing core fibre network capacity. Using this innovative approach could enable BT to bette[illegible]ds of consumers and busi[illegible]es [illegible]g ever more bandwidth-hungry services, such as video streaming.

Cormac Whelan, head of sales for Alcatel-Lucent [illegible] & Ireland said: "As part of our [illegible]ng relationship, BT and Alcatel-[illegible]tinue to work together to use innovation from Bell Labs, Alcatel-Lucent and BT Research & Innovation to move the industry forward and meet the ever [illegible] the marketplace. These [illegible]her step forward by [illegible]nt in this continual evolution."

INVESTING IN THE FUTURE

Th[illegible]rough gives BT the opportunity [illegible] new features and technology without the need to update the existing optical transport infrastructure.

Neil J. McRae, chief network architect at BT said: "Investing for the future is core to BT's strategy and this outstanding achievement demonstrated that BT can easily introduce new features and technologies across our core network maximising the efficiency of our existing infrastructure. Working with Alcatel-Lucent on this trial has been highly productive in demonstrating the viability of an alien wavelength approach."

Originally published January 28, 2014

Factors driving PSM4 and silicon photonics for data center architectures

by **BRENT HATFIELD,** Molex Inc.

MULTIPLE FORCES AND trends are driving the implementation of parallel single-mode quad (PSM4) and silicon photonics (SiPh) in next-generation data center designs and architectures. PSM4 and SiPh fill an important gap in data center interconnect options – the gap between the reach of 25G multimode options and that of long range (LR) optical approaches. Using pigtail designs with these technologies can ease implementation via easy mating to structured cabling. From a "future proofing" perspective, PSM4 is an approach that can accommodate both current and future bandwidth upgrades.

The cloud

While most consumers have used the cloud in one form or another for many years (Flicker, iTunes, etc.), one recent major trend has been the aggressive proliferation of cloud use by private enterprises. A 2011 study by IBM found that 70% of mid-sized businesses were using cloud-based analytics and 66% had either already deployed or were planning to deploy cloud-based technologies.

Rather than increase IT capacity internally, moving items to the cloud has proven to be both cost-effective and timely. Microsoft recently reported that its Skydrive service has 17 million customers that store 10 petabytes of data. The trend continues to accelerate.

Entertainment on-demand and entertainment mobility

From Hulu to Apple TV to Amazon Prime, Netflix, and others, entertainment is transitioning from the set-time model to an on-demand model. Furthermore, entertainment is now viewed as a product that is expected to be available not just when you want it but also where you want it.

Multiple viewing/experience platforms and options exist. Internet-connected televisions, tablets, PCs, and smartphones enable viewers to access any level of information or entertainment at the place and time of their choosing. This, in turn, requires large-scale infrastructure to store and support programming available on-demand by consumers.

Big Data

According to Roberto V. Zicari of Object Database Management Systems, "Every day, 2.5 quintillion bytes of data are created. This data comes from digital pictures, videos, posts to social media sites, intelligent sensors, purchase transaction records, cell phone GPS signals to name a few. This is Big Data."

Previously most of the data was released and deleted, viewed as having little value. Or, at a minimum, it was viewed as too cumbersome to store and use. However, recent new technologies in database management (most notably the development of Hadoop) enable large-scale data-based relational information, even where the data is in disparate types, formats, and across different systems. Essentially, all data has become both manageable and potentially useful. As a result, all data has begun to be stored indefinitely and mined for information and relationships.

New data center architectures

As a result of these multiple factors, data centers are becoming larger and more distributed. Virtualization means that system architectures that once took additional hardware can be built with less equipment, but with more interconnections for ensuring resources are fully used and interactive. Next-generation data center designs typically require two main attributes: scalability and uniformity of performance coupled with low latency. Scalability is required to handle the expansion of services, addition of customers, and the increase in data. Uniformity of performance is required to provide a smooth flow of data between nodes.

The traditional three-tier data center architecture of core, aggregation, and access was well suited for traditional use of email, webpages, and traffic. However, new requirements for video and content delivery, virtual machines, cloud access, and social networking content assembly required low latency across the data center.

Non-blocking bandwidth became both a requirement and a technical challenge that demanded a new data center architecture.

"Spine and Leaf" architectures resolve these two issues. A central spine handles high-bandwidth data between leafs, whereas a leaf controls the traffic flow between clusters of servers. The performance is thus balanced, while the structure supports the addition of leafs as the system scales. Traffic between nodes is balanced and accessible with low-latency east-west traffic flow.

This type of architecture has many attributes; one physical one is the need for longer links. And that could mean a significant change in cabling.

Multimode fiber and the distance/data rate challenge

In the last decade, copper high-speed interconnects have faced a primary challenge. As data rates increased, the distances that copper interconnects could accommodate decreased to the point that alternative approaches had to be found. To handle increased data rates, new and exotic ways were developed to expand the reach of copper: higher awg cable, new dielectric materials, and new equalization schemes, CDRs, and other active signal integrity devices. Each method came with increasing costs and implementation challenges. Ultimately, the crossover point occurred as multimode fiber became the only viable solution for exceeding a certain combination of data rate and distance.

Over time, however, the same challenges have arisen with multimode optics, propelling demand for more exotic fiber and connector constructions:

- 1GBase-SR with OM2–fiber reach is 550 m
- 40GBase-SR4 with OM3 – fiber reach is 150 m
- 100GBase-SR4 with OM4–fiber reach is 100 m.

Each new upgrade requires significant data center down time, including additional price premiums for structured cabling. And the time between these upgrades is shrinking. 10 Gigabit Ethernet (GbE) is deployed, and 40GbE systems are being installed now. We'll soon see 100GbE begin deployment, and a consortium of suppliers announced an MSA for 400GbE last year. Terabit Ethernet is on the horizon. The traditional economic model of replacing an outdated infrastructure and expecting pay off within a few short years is increasingly under pressure.

Bandwidth and technical challenges have been aptly referred to by some as a “coming data tsunami.” The other significant challenge is rising cost – and consumer expectations. Rightly or wrongly, consumers have come to expect noticeable performance increases with little or no cost increase. From photo transfer, to video, to HD video, to “everything available on demand through every device,” customers want more but do not want to pay more. This places increased pressure on providers to lower costs.

At a time of increasing cost pressures, constant multimode cabling upgrades run counter to industry needs.

A solution: SiPh PSM4 for 4-km reach

A SiPh-based PSM4 approach helps resolve the multiple challenges of needing higher bandwidth, longer distances, low power, and future proofing. SiPh-based long-reach PSM4 products offer seven primary advantages:

1. **Distance:** With transmission distances up to 4 km, they can accommodate most, if not all, new data center requirements.
2. **Power:** SiPh-based active optical cables (AOCs) use basically the same power consumption as VCSEL-based products: under 1 W for a 10G QSFP+ AOC and 1.5 W at 25 Gbps.
3. **Cost:** Cost is roughly the same as VCSEL-based AOCs, but as they are singlemode they use much less expensive singlemode cable as a transmission medium. As VCSEL-based product speeds increase, they require ever more expensive types of fiber to transmit effectively. After most networks have upgraded to OM3, another upgrade to OM4 has serious cost implications. And what comes after OM4, and how soon will it be needed? With singlemode SiPh AOCs the fiber stays inexpensive singlemode – and consistent as data rates increase.
4. **Future-proof structured cabling:** This is becoming an increasingly important factor in driving the adoption of SiPh in the market, particularly in new installations. Upgrading structured cabling is less economically viable as speeds continue to increase. PSM4 can accommodate:

- 10GbE – 4x10 Gbps
- 40GbE – 4x10 Gbps aggregated
- QDR Infiniband – 4x10 Gbps
- FDR Infiniband – 4x14 Gbps
- EDR Infiniband – 4x25 Gbps
- 100GbE – 4x25 Gbps aggregated.

5. **Pigtail plug-in options for ease of installation and upgrades:** PSM4 AOCs are available in pigtail versions. With a pigtail, one end is a traditional QSFP+ active interface, and the other either MPOs or LCs. This enables a quick connection to structured cabling (so the AOC side doesn't have to be pulled long distances) and the ability to quickly upgrade to new products as needed. So when the 4x28G zQSFP product is introduced, it can be immediately connected to structured cabling already in place. Upgrades from 10GbE/40GbE to 100GbE can be done immediately with no costly structured cabling upgrade and associated data center downtime.
6. **SiPh actives becoming broadly available:** An ecosystem of proven, reliable, cost-effective SiPh devices exists to accommodate singlemode data center links. SiPh is widely deployed today in PSM4 AOCs, and several new entrants into the market are offering a variety of options.
7. **SiPh devices offer a clear technology path beyond 25G:** VCSELs face increasing design challenges as speeds increase. On the other hand, with SiPh-based systems most modulating schemes have a clear and well-understood technology path to 50G, 100G, and beyond.

The future is now

Until recently, it was generally believed that bandwidth/distance design challenges were limited to high speed copper interconnects; most thought the transition to fiber optics in the nineties resolved these issues. Few predicted that traditional VCSEL-based optics would begin to experience some of the same challenges as copper interconnects did over a decade ago. But they have.

Singlemode SiPh bridges the gaps caused by the dual requirements of longer distances and higher data transmission speeds. For many customers this technology provides a lower-cost, lower-power option for what can be referred to as medium-reach distances that is future-proof for the next generations of data transmission speeds. With a broad base of products both available and being developed by multiple suppliers, PSM4 and SiPh AOCs can be deployed today to enable new data center architectures, with the assurance that they'll provide a path to cost-effectively meet future generation upgrade requirements.

BRENT HATFIELD is product manager, Fiber Optics Division, at Molex Inc.

Alcatel·Lucent

Company Description

Alcatel-Lucent is at the forefront of global communications, providing products and innovations in IP and cloud networking, as well as ultra-broadband fixed and wireless access, to service providers, enterprises and institutions throughout the world. Underpinning Alcatel-Lucent in driving transformation across the industry is Bell Labs, an integral part of Alcatel-Lucent and one of the world's foremost technology research institutes.

EXTERNAL RESOURCES

- Podcast: Bring Agility to Your Optical Network
- Lightwave OFC 2014 Podcast Series
- Technical White Paper
- VIDEO: More than moving light
- Lightwave Podcast Series 2013