

Breaking the Bandwidth Bottleneck



Today's carriers support more than 400 million global broadband subscribers worldwide. This subscriber base and its bandwidth appetite have grown so quickly that industry observers expect bandwidth demand to outstrip availability within just a few years. Monthly traffic, worldwide, is growing at a rate of up to 60 percent, annually, according to the University of Minnesota, and traffic volume reached 5000 to 8000 petabytes (PB) by the end of 2008. In carrier networks, the fiber plant uses Dense Wavelength-Division Multiplexing (DWDM) technology to combine more than 80 channels to transport the Internet traffic. Today, these DWDM channels are dominated by 10Gb/s line rate and channels are separated by a 50GHz spacing. With limited fiber, these 10G channels are the constraints at the critical fiber route now.

The problem isn't restricted to carrier networks - the enterprise is also devouring bandwidth for software-as-a-service (SaAs), remote backup and other applications that are hosted on vendors' servers and require high-bandwidth Internet transmission. As enterprise bandwidth appetite grows, aggregation becomes a bottleneck that, in turn, feeds directly into an increasingly bandwidth-constrained core.

100G is the answer, but getting there is no small task. We need a faster path to 100G than the ten-year path from 10G to 40G (beginning with standards development and moving through field trials to this year's early deployments). 40G got an early start, which helped to accelerate adoption of bandwidth-consuming applications, but then the telecommunications downturn hit, slashing R&D investments for important sections of the 40/100G component supply chain just when they were needed most. Now we're in catch-up mode, and the industry is feverishly finalizing 100G standards and modulation techniques while resolving key challenges including chromatic and polarization-mode dispersion.

The Institute of Electric and Electronic Engineers (IEEE), International Telecommunications Union (ITU) and Optical Internetworking Forum (OIF) have been working on standards since last year, with only minor changes expected prior to finalization in 2010. This will enable further 40G deployment in carrier network and

Breaking the Bandwidth Bottleneck

Published on Electronic Component News (<http://www.ecnmag.com>)

100G field trials in 2009, followed by mainstream 40G adoption and 100G deployments the following year.

In IEEE802.3ba frame work, the 100G Ethernet is a parallel 10x10.3G physical interface (CAUI) from the 100GE PCS. For the inter-office application, the polarization mode dispersion (PMD) specifies a 4x25.8Gb/s optical interface, which is optically multiplexed into a single mode fiber over a distance beyond 10km. It requires a bit-multiplexer to generate 4x25.8G/b bit streams from the 10x10.3G CAUI signal.

The ITU has ratified schemes to map 40G/100G Ethernet payloads into OTU3/OTU4. The data rate of the optical transport unit (OTU) that is carrying 100GE (OTU4) and the forward error correction (FEC) overhead has also been specified. The ITU has also considered transmission by different modulation formats so that extended link budget can be met. One example is the use of dual-polarization differential quadrature phase shift keying (DP-DQPSK) modulation. With this configuration, a single wavelength optical carrier carries 4-bit DQPSK precoded signals at a symbol rate of 28Gbaud/s. Since DP-DQPSK operates at a symbol rate that is one-quarter that of its original serial data rate, the tolerances to PMD, chromatic dispersion (CD) and DWDM filters are significantly improved.

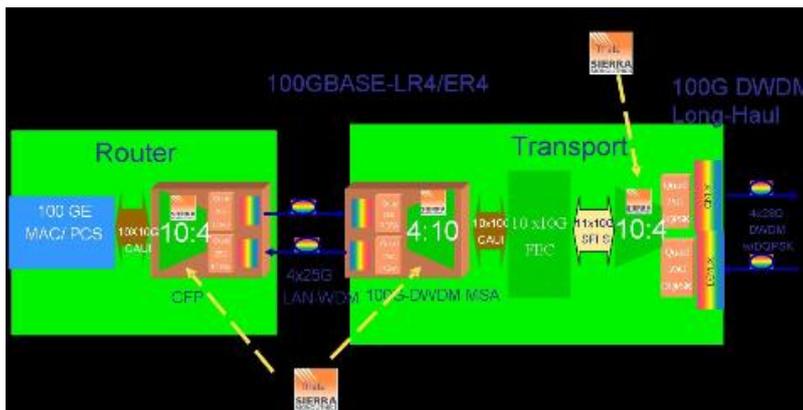


Figure 1 is a simplified block diagram showing how a 100G Ethernet signal from media access controller and physical layer solution (MAC/PCS) is transported to/from a central office and through the optical network.

The Optical Internetworking Forum (OIF) 100G Ultra-Long-Haul DWDM Project has selected the dual-polarization QPSK (DP-QPSK) as its modulation format for recommended implementations. The use of coherent detection approach is also specified as a breakthrough advance beyond traditional optical transmission technology. The consensus is that traditional direct detection techniques may not be enough to provide the necessary transmission performance in today's 10G network. On the other hand, coherent DP-QPSK offers the promise of providing superior performance beyond today's 10G DWDM benchmark.

None of this can be accomplished, though, without a complete 100G ecosystem. Optical transponders require many components: multiplexer/clock multiplier unit (CMU); driver modulator and modulator/laser; demodulator (photodiode/transimpedance amplifier); and clock data recovery (CDR)/demultiplexer. Fortunately, with breakthroughs in semiconductor design

Breaking the Bandwidth Bottleneck

Published on Electronic Component News (<http://www.ecnmag.com>)

technology and evolutionary process technology advances, pioneering 40G components have been demonstrated and adopted in telecommunications applications and next-generation 100G Ethernet solutions are well positioned for mass production and large-scale deployments.

100G Mux/Demux Advances

As shown in Figure 1, a 10:4 multiplexer and demultiplexer are needed for both Ethernet and DWDM transceivers. Some distinctive features are required for each application. For example, the electric interface protocol for 100G Ethernet is CAUI, while the SFI-S interface is more desirable for telecommunications applications. For Ethernet, the signal is always terminated at the receiver. The long-haul DWDM is typically relayed by multiple spans of single-mode fiber. Each span features an Erbium-doped fiber amplifier (EDFA) and dispersion compensation fiber, and there may also be a repeater in the span. Therefore, a more stringent jitter performance is required for long-haul telecommunications applications. This also results in different implementations at Ethernet PCS and OTU framer/FEC. For example, the Ethernet host board typically has a pluggable connector and long printed circuit board (PCB) traces to host multiple transceivers, whereas the interconnect in telecommunications host board is typically surface-mounted. Therefore, pre-emphasis and equalization is needed for the Ethernet application.

One example of how 100G Mux/Demux solutions can be architected is Sierra Monolithics' SMI10021 and SMI10031 devices - together, the industry's first multiple-rate 100G Mux/Demux chipset. It is suitable for both Ethernet and telecommunications applications. On the Mux side, an integrated DQPSK pre-coder features user-selection for dual-polarization or single-polarization pre-coding for DQPSK applications, and also supports LAN-WDM applications, in which the pre-coder is disabled. A jitter clean-up phase-locked loop (PLL) option is available for meeting stringent jitter performance requirements. A pre-emphasis option is provided for the Mux outputs, which can be used to compensate for interconnect loss through the printed circuit board (PCB). An on-chip pseudo random bit sequence (PRBS) generator and error checker are an integral part of the chip, with industry-standard PRBS patterns to facilitate bit error testing for the chip, transceiver and system. Optional equalizers are also included for the client-side interface, which supports both CAUI and SFI-S protocols. Figure 2 shows a block diagram of Sierra Monolithics' SMI10021 10:4 MUX/CMU.

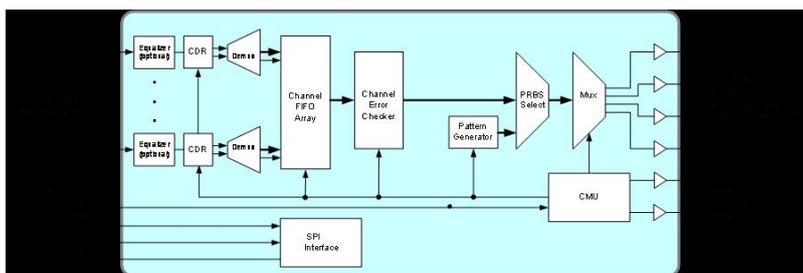


Figure 2: Theta-100G SMI10021

10:4 MUX/

To execute all necessary Demux tasks, an integrated DQPSK decoder is also included for test purposes. Clock and data recovery units (CDRs) are positioned on

Breaking the Bandwidth Bottleneck

Published on Electronic Component News (<http://www.ecnmag.com>)

each input data channel. Optional decision threshold and data sampling phase adjustment are provided for each input data line, so that sensitivity can be optimized depending on input conditions. Similar to the Mux, the Demux also contains a PRBS generator and error checker for bit error test. Pre-emphasis for the client-side interface is also an included option to compensate for interconnect loss in PCB. Figure 3 shows the functional blocks in the Sierra Monolithics SMI10031 4:10 CDR/DEMUX.

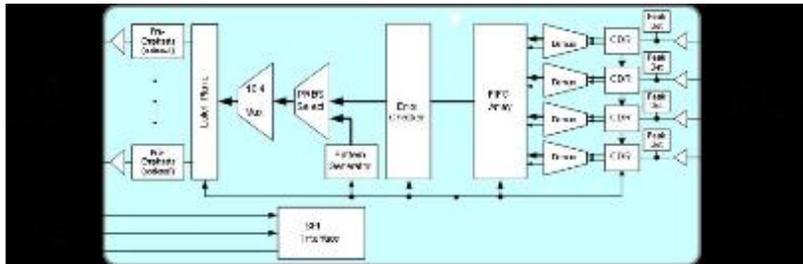


Fig. 3: Sierra Monolithics'

SMI10031 4:10 CDR/DEMUX device

Rounding out this Mux/CMU and Demux/CDR architecture are a 10x10.3Gb/s (MLD/CAUI) or 11x11.2Gb/s SFI-S interface on the client side, as well as a de-skew function in compliance with OIF SFI-S, plus a line side pre-skew. The DQPSK pre-coding function is implemented with dual I/Q-interleaved outputs (4x28Gbps) for DP-DQPSK applications, and may also be configured to enable a single-polarization 2x56Gb/s DQPSK modulation with a pair of external 2:1 Mux/Demux through use of the synchronous high-speed clocks. Typical jitter swing is 3.7psec p-p typical, and differential output level is 0.6 to 1.2V p-p.

Process technology is another critical piece of the 100G puzzle. Bipolar complementary metal-oxide semiconductor (BiCMOS) silicon germanium (SiGe) technology has been battle-tested in 40G deployments and also provides the foundation for next-generation 100G solutions. The f_T value for a bipolar block is 210GHz, which provides a lower noise floor and wide dynamic range that is optimal for 40/100G device's eye quality requirements. The 130nm CMOS block is suitable for 10Gb/s application. BiCMOS SiGe technology improves performance at lower power while enabling higher levels of integration which translates into superior performance, compact footprint and low cost for optical transceivers.

Expect several key developments during 2009. As the industry solidifies 100G standards, component and system vendors will start releasing solutions. A variety of transponder components will become available, enabling early deployment and system trials.

100G's benefits are enticing, including bandwidth relief for over-taxed networks, and cost-per-bit reduction for the existing fiber infrastructure with significantly improved network economics. 100G standards are also expected to unite optical transport with Internet Protocol to deliver a simplified network and minimize traditional payload mapping requirements.

Breaking the Bandwidth Bottleneck

Published on Electronic Component News (<http://www.ecnmag.com>)

Taqi Mohiuddin is Director of Marketing for the 40G/100G Optical SerDes product lines at Sierra Monolithics since September 2008. He has a B.S. in Electrical Engineering from University of Illinois and an MBA from DePaul University. He was formerly VP of Marketing for Siverge Networks, and Executive Director of Marketing for Mindspeed Technologies.

Email: Taqi.Mohiuddin@monolithics.com [1]

Kun-Wook Chung is a Senior Applications Engineer for the 40G/100G SerDes products at Sierra Monolithics. He was formerly design engineer in optical amplifier group at Agere systems and senior design engineer at Circadiant Systems. He received his Ph.D degree in electrical engineering from Pennsylvania State University, University Park in 2001. From 1992 to 1997, he was a research engineer in Satellite Communication division at Electronics and Telecommunications Research Institute, Korea. He has published more than 30 journal and conference papers in fiber optics, Microwave circuit design.

Email: KChung@monolithics.com [2]

Song Quan Shang is the systems architect for optical products at Sierra Monolithics since March 2008. He leads the product definition, architect and specification for high-speed Serdes products at 40Gb/s and 100Gb/s, which has an application in telecommunication, Ethernet and server interconnect. He is also responsible for the industry standard developments at IEEE802.3ba and OIF. Prior to SMI, Song was a principal engineer at Intel Corporation between 2002 to 2008. He was with Ciena Corporation developed long-haul and metro DWDM system from 1998 - 2001.

Email: Song.Shang@monolithics.com [3]

Source URL (retrieved on 07/10/2014 - 9:48pm):

http://www.ecnmag.com/product-releases/2009/05/breaking-bandwidth-bottleneck?qt-video_of_the_day=0

Links:

[1] <mailto:Taqi.Mohiuddin@monolithics.com>

[2] <mailto:KChung@monolithics.com>

[3] <mailto:Song.Shang@monolithics.com>