

Electronic dispersion compensation brings native 10 Gbps to networks

AT 10 GBPS, DISPERSION HAS A DOMINANT EFFECT ON OPTICAL-LINK PERFORMANCE FOR LONG- AND SHORT-HAUL-NETWORKING APPLICATIONS. TO HANDLE THIS DISPERSION, DEVELOPERS MUST EITHER UPGRADE THE FIBER INFRASTRUCTURE OR IMPLEMENT DISPERSION COMPENSATION.

Bringing new efficiencies to high-speed-optical-communications applications continues to challenge scientists and engineers alike. Whenever they overcome the current set of challenges, the demand for higher bandwidth over greater distances pushes the limitations of technology. Effects that designers previously ignored because they were below the noise floor at lower data rates and distances suddenly become significant design barriers.

Today, 10-Gbps, long-haul, and metropolitan SONET (synchronous-optical-network) OC (Optical Carrier)-192 optical links over SMF (single-mode fiber) can reach only to approximately 80 km, primarily because of impairments in the fiber. Similar problems arise in data-center and backbone applications in which 10-Gbps Ethernet links operating over legacy OM1 (optical-module) MMF (multimode fiber) can reach no more than 26m because of the effects of signal dispersion at these higher data rates. What was acceptable noise at 1 and 2.488 Gbps is debilitating at native 10 Gbps.

IT managers and carriers need to be able to cost-effectively scale networks using the infrastructure. At OC-48 and 1-Gbps rates, links can operate beyond 80 km and 26m distances with little dispersion and without significant signal-integrity impact. At 10 Gbps, however, dispersion has a dominant effect on optical-link performance for most long- and short-haul-networking applications. To handle this dispersion, developers must either upgrade the fiber infrastructure or implement some form of dispersion compensation.

Currently, both the OIF (Optical Internetworking Forum) and the IEEE have taken on the challenge of increasing the reach and reliability of 10-Gbps technology. They will address dispersion for both long- and short-haul applications with EDC (electronic-dispersion-compensation) technology to account for optical dispersion in the electrical domain. Through the use of these standards, designers can consistently apply EDC throughout networks, enabling carriers and IT managers to cost-effectively and reliably upgrade their networks.

In collaboration with the ITU (International Telecommunications Union), the OIF is defining the SMF EDC Standard under the ITU-TSG15. The standard defines long-reach, 10-

Gbps, OC-192 SONET links operating over 145-km distances or 120 km with worst-case fiber. It enables seamless upgrades from OC-48 without the need to replace fiber or deploy expensive and bulky DCF (dispersion-compensated fiber). The organizations intend that the standard will address longer reach applications in which the minimum chromatic dispersion must be at least 2400 psec/nm.

For short-reach applications, the IEEE is developing the 802.3aq standard for upgrading 1-Gbps links to 10 Gbps over MMF. Most applications now run 1-Gbps links, and a few 10-Gbps deployments use parallel formats of 10G-BaseLX (local-exchange) 4 PMD (polarization-mode dispersion)—that is, four 2.5-Gbps links over OM1 fiber at distances to 300m. The 802.3aq standard will run serial 10-Gbps links over legacy OM1 MMF at distances to 220m.

Beyond bandwidth and distance considerations, the market factors driving the development and deployment of these two standards include legacy-fiber upgrades, lower costs, and increased overall link reliability. In the case of long-reach ITU-TSG15, carriers will be able to replace transponder modules, along with appropriate back-end components, such as framers, with 10-Gbps transponders. As a result, they will be able to upgrade equipment without upgrading the fiber.

One of the primary attractions of short-reach 802.3aq is that




Figure 1 Because each light wavelength comprises different colors, the greater the distance light travels, the wider the pulse of light that spreads. Chromatic dispersion results in adjacent pulses that interfere with each other, leading to intersymbol interference.

carriers can run 10-Gbps links in a native serial format using a less complicated module than LX4. LX4 modules use four wavelength-stable lasers and a complex optical multiplexer, which increase overall system cost and require detailed integration and testing. In comparison, 802.3aq modules require only one wavelength of light. As a result, they consume less power, are easier to maintain and are thus more reliable, and are approximately one-half the cost of 10G-BaseLX4 PMD links. Additionally, manufacturers promise even lower cost and smaller modules that they will base on the XFP (10-Gbps small-form-factor pluggable module).

LX4 PMD links are also more reliable than current links. Manufacturers are now developing or shipping products employing these standards, effectively compensating for known sources of interference and effecting substantial improvements in signal quality and therefore overall link reliability (Table 1).

SOURCES OF DISPERSION

EDC addresses three main types of interference that lead to link impairment at 10-Gbps rates: chromatic dispersion, modal dispersion, and PMD. These types of interference depend on the symbol rate—that is, as signal speeds increase, dispersion has a greater effect. Hence, these types of dispersion have taken center stage in the migration to 10-Gbps networks.

Chromatic dispersion, the result of physical and waveguide properties, manifests itself as the spreading of a pulse of light as it travels over great distances. Optical lasers output pulses of light with a finite spectrum comprising colors. The longer the fiber over which the pulse travels, the wider the pulse spreads out (Figure 1). Difficulties arise when the resulting energy from a pulse begins to interfere with that of an adjacent pulse. This interference causes ISI (intersymbol interference) in the electrical domain. The spreading of symbols across each other causes errors; the receiver side of the link cannot easily distinguish the symbols because they are no longer at ideal levels. Depending on the fiber, pulse spreading may cross several UIs (unit intervals); a dispersion of one UI means that adjacent symbols within the same symbol string begin to interfere with each other.

SMFs typically have a dispersion slope of about 17 psec/nm at 1550 nm, or approximately the operating range of a long-haul transmission system. Manufacturers quantify chromatic dispersion by the distance light travels along the fiber. A pulse with a center frequency of 1550 nm transmitting over 140 km would experience a total chromatic dispersion of approximately 2400 psec/nm, which is equal to the specification for the OIF's upcoming ITU SMF long-reach standard.

Interference between modes of light arriving at a receiver at different times causes modal dispersion. As a result, modal dispersion is specific to MMF in short-reach data centers and backbones. Modal dispersion arises from imperfections in fiber that progressively degrade light, causing the light to spread, disperse, and eventually overlap (Figure 2).

PMD, typically a concern of SMF applications, is a phenomenon in which a single pulse appears as multiple pulses far-




Figure 2 Imperfections in fiber and continued degradation over time cause modes to propagate through the fiber at different speeds. This situation leads to dispersion and potential overlapping of light.




Figure 3 PMD causes a single pulse to spread at the far end because of the delayed arrival of the two perpendicular polarization modes within the fiber.

ther down the fiber (Figure 3). Optical fiber supports two perpendicular polarization planes, and ideal fiber would transport both polarization signals to arrive at the receiver side at the same time, appearing as a single pulse. However, fiber is neither perfectly round nor stress-free, which leads to phase shifting of the pulse. Designers can compensate for PMD using standard receivers for applications requiring reaches of less than 80 km. As link distances increase, however, the effects of PMD are statistical and complex to measure. If compression or kinks have damaged the fiber, for example, performance degrades appreciably more quickly; consider that a single kink could cause two components of light to travel at 90° to each other. For this reason, the condition of the fiber can have a more pronounced effect on signal integrity than the length of the link.

COMPENSATION OPTIONS

You can effectively implement EDC using a variety of equalization algorithms. The three most common are CTFs (continuous-time filters), FFE/DFE (feedforward-equalizer/decision-feedback-equalizer)-algorithm combinations, and sophisticated MLSE (maximum-likelihood-sequence-estimator) equalization. CTFs offer the simplest, most cost-effective, and lowest power EDC implementation. By boosting or bandlimiting the signal within the frequency band of interest, a CTF can adjust the analog bandwidth of the optical front end, effectively acting as a lowpass filter. By amplifying certain frequencies and attenuating others through waveshaping, a CTF can compensate for chromatic dispersion. However, you can reduce high-frequency noise only so much before the CTF begins to filter the signal, as well, severely curtailing such compensation. Thus,

CTFs are appropriate only in applications in which dispersion is not excessive.

FFE/DFE algorithms apply a more sophisticated approach to compensation than that of a CTF. FFE/DFE implementations use multitap algorithms to compensate for ISI that exceeds one UI of interference. An EDC implementation comprises an AGC (automatic-gain-control) block, a CTF/FFE block, a DFE block, a CDR (clock- and data-recovery) block, and an LMS (least-mean-squared) adaptation block (Figure 4).

When there is only a single UI of interference, compensation involves determining whether an adjacent symbol has spread into the current symbol and then adding or subtracting the symbol. When more than one UI of interference is present, a symbol can spread and distort several adjacent symbols, making compensation more complex. FFE removes distortion before a symbol's primary energy point or precursor area. DFE compensates for interference following a symbol's primary energy point or postcursor area.

MLSE implementations provide even more sophisticated equalization architectures. Incorporating Viterbi-decoder algorithms, an MLSE requires a DSP approach to filtering. Although an MLSE can achieve better performance than an FFE/DFE, DSP implementations are generally more complex and often consume two to four times the power. For these reasons, MLSE-based approaches most often find use in applications having little room for compromise in performance. For example, MLSE compensates for severe nonlinearity in fiber or in ultra-long-haul applications.

The most common EDC implementation uses a combination of FFE and DFE, providing a higher level of performance and reliability than do CTFs but at a more reasonable power cost than MLSE implementations require. In addition, an analog FFE/DFE design typically has lower power dissipation than a digital implementation because there is no need to convert the analog signal into the digital domain using high-speed ADC or DSPs. Typically, designers base an FFE implementation on an analog distributed amplifier using various on-chip transmission lines to create delay elements. The DFE portion uses sample data to determine signal quality and requires a bit-rate clock, so you can implement it primarily in analog or primarily in digital, depending on your application's architecture.

Power consumption is only one consideration, however. You must also consider performance stability over extreme operating conditions, as well as other issues. For example, model dispersion in MMF is often more of a factor than it is in SMF. As




Figure 4 An EDC implementation comprises an AGC block, a CTF/FFE block, a DFE block, a CDR block, and an LMS (least-mean-squared) adaptation block.

a consequence, equalization for short-reach MMF is often more sophisticated than that for long-haul SMF. EDC also has a substantial effect on signal reliability. Traditional receivers without EDC can recover an optical signal only if the dispersion is less than approximately one-half UI over the length of the fiber. The new IEEE 10-Gbps standard, however, supports runs as long as 220m using OM1-type 62.5-micron fiber and specifies that the receiver must be able to handle more than four UI of dispersion. Without EDC, you cannot possibly meet this requirement.

DYNAMIC EDC

Adaptability is an essential characteristic for any successful network. Every optical link has characteristics that its length, quality, condition, and other important factors determine. Equipment needs to automatically adapt to a link when the user installs it to achieve the best performance, efficiency, and reliability. Additionally, as fiber degrades over time and introduces new sources of interference, such as new kinks in the fiber, line cards must refine compensation algorithms to adapt to these changes. In this way, designers can achieve further cost savings by eliminating the need for hand-tuning links for reach and wavelength. Rather, with self-adaptable EDC-enabled equipment, users can install line cards without manual tuning for true plug-and-play deployment. An adaptive EDC implementation can improve more than just reliability. Because designers can tailor equalization for an application, adaptive EDC also facilitates the use of a single-board design across multiple applications.

Self-adaptation requires closed-loop-feedback mechanisms that enable equipment to calibrate itself by slightly modifying filters and gains that improve signal response until the system achieves an ideal signal. Using well-established LMS algorithms to implement EDC, you can easily implement efficient self-adaptation. Many silicon vendors are looking to integrate EDC directly on transceivers to further simplify this process for developers.

Clearly, 10-Gbps Ethernet is an important enabling technology in

TABLE 1 SIGNAL QUALITY WITH AND WITHOUT EDC

Technology	10-Gbps Ethernet short reach	10-Gbps Ethernet short reach, with EDC	10-Gbps SONET long haul	10-Gbps SONET long haul, with EDC
Standard	802.3ae	802.3aq	G.959.1	G.959.1
Wavelength (nm)	850	1300	1550	1550
Fiber (microns)	MMF (62.5/125)	MMF (62.5/125)	SMF (9)	SMF (9)
Maximum distance	26m	220m	80 km	145 km
Dispersion	NA	NA	1600 psec/nm	2400 psec/nm
Bit-error rate (Hz)*	1×10^{-12}	1×10^{-12}	1×10^{-12}	1×10^{-12}
Path penalty*	NA	NA	2 dB	2 dB*

*Without EDC, the path penalty would be more than 10 dB, and bit-error rate would be less than 1×10^{-6} .

data-center, storage-network, and back-haul applications. Current 10-Gbps implementations have aggregated only four 2.5-Gbps links because of the effects of signal dispersion at native 10 Gbps. With the means to compensate for dispersion at high signal frequencies, IT managers and carriers will be able to cost-effectively upgrade 1-Gbps and OC-48 links to native 10 Gbps, all without laying new fiber or deploying bulky DCFs.

EDC is an effective means of compensating for dispersion and does so across the optical spectrum from short- to long-haul applications. By compensating for optical dispersion in the electrical domain, EDC enables developers to upgrade network links without changing optical components. It also significantly

increases overall reliability and the distance that links can run.

EDC is an essential ingredient for the successful deployment of native 10-Gbps links. Both the OIF and the IEEE recognize this fact and have been managing the development of EDC to ensure that it stabilizes, rather than—lacking an industry

standard—impedes, the 10-Gbps-Ethernet market. As expected, the OIF's ITU-TSG15 is nearing ratification. Interoperability testing between industry leaders is under way, and the organization should approve the standard with no fundamental changes. Likewise, the IEEE's 802.3aq standard is making significant headway. Currently in draft status, the standard should achieve ratification by midyear. In the meantime, developers

can continue to design 10-Gbps equipment knowing that they can achieve the performance, distance, cost, and reliability expectations of the market. **EDN**

AUTHORS' BIOGRAPHIES



Michael Furlong is a senior product-line manager with Broadcom Corp. He manages the business and strategic direction of the company's high-speed-interconnect-product line. He received a bachelor's degree in electrical engineering and a master's in business administration from the Florida Institute of Technology (Melbourne).



Ali Ghiasi, PhD, is the chief architect for Broadcom Corp's optical business. His current interests include EDC (electronic-dispersion-compensation) applications for datacom and telecom, storage networking, and next-generation optical interfaces, and he is a founder of the XFP (10-Gbps small-form-factor pluggable module) MSA (multisource agreement). Ghiasi has a doctorate in electrical engineering from the University of Minnesota (Twin Cities) and master's and bachelor's degrees from North Dakota State University (Fargo). Ghiasi has written for more than 50 publications, as well as architectural proposals for the standards bodies IEEE, OIF, FC, and InfiniBand. In addition, he is the named inventor on more than a dozen patents and patent applications in high-speed communications and interconnects.

MORE AT EDN.COM ▶

➕ Go to www.edn.com/ms4182 and click on **Feedback Loop** to post a comment on this article.