Open Column

by Anthony Chan Carusone

The Limits of Light: The Finite Bandwidth of Optical Fibre

Optical fibre is regularly touted as the ultimate broadband link—an "infinite"-bandwidth medium towards which all wireline communication technologies must inexorably migrate. After all, photons are faster and better behaved than electrons - right? The reality is somewhat different… Dispersion has limited our ability to communicate using light. This article summarizes the open problems facing researchers that strive to extend those limits. For example, there has been significant recent progress on linear methods of dispersion compensation using integrated electronics, but further work is needed to lower the power and area of those methods. Decision-feedback compensation offers better performance but requires even higher power consumption. To realize the full potential of optical fibre links, researchers are looking further ahead towards maximum likelihood sequence estimation, orthogonal frequency-division multiplexing (OFDM), and/or iterative error correction. However these methods demand digital processing thus opening problems in very high-speed analog-to-digital and digital-to-analog conversion. Research on adaptive optics opens the exciting possibility of diversity communication for optical links operating roughly 100$\times$ faster than existing wireless diversity transceivers. However, rapid time-variations in dispersive channels are particularly difficult to track using optics. In summary, dispersion compensation currently faces a rich variety of open problems in circuits and systems.

Of course cost is king, so to have an impact the solutions must be economical. Electrical energy simply can not be efficiently converted into optical energy and back again using our least expensive integrated circuit technologies. Furthermore, optical components are notoriously finicky and expensive. Hence, optical links are only cheaper than their electrical counterparts when these disadvantages are offset by simpler, lower power transceivers in smaller and less expensive packages. This seems reasonable since using optical fibre obviates the need to combat nasty electrical channel impairments such as skin effect and dielectric loss. But in practice, it is often difficult to realize these advantages for two reasons.

Digital Object Identifier 10.1109/MCAS.2008.923059

First, remarkably low-cost methods for compensating electrical link impairments have been developed. Consider chip-to-chip communication links. I recall a cover story from the August 2002 IEEE Spectrum promising that consumers would begin to see optical chip-to-chip links in their computers within 5 to 10 years. But here we are, 6 years later, and electrical chip-to-chip links continue to outperform their optical counterparts. Why? Because considerable research effort has resulted in electrical chip-to-chip transceivers occupying less than 0.5 mm^2, consuming less than 2.5 pJ/bit, and completely integrated in CMOS [1]. Optical technologies simply can not compete with numbers like that.

The second reason is that optical links also suffer from channel impairments and, therefore, may require sophisticated signal processing of their own to support robust communication. Light (whether or not it is confined to an optical fibre) is dispersive. Since different portions of an optical pulse propagate with different velocities closely-spaced digital pulses interfere, just as in electrical links. The result is, effectively, a bandwidth limitation in our infinite-bandwidth medium! Of course these bandwidth limitations can be mitigated with sufficient signal processing, but that means more sophisticated transceivers consuming more power. Hence, there is considerable interest in both optical and electronic dispersion compensation (EDC), the focus of this article.

But before describing the signal processing employed to combat dispersion in modern optical links, a little background on dispersion in optical fibres is required. This is provided in Section 1. Then, Sections 2, 3 ,4, 5, and 6 describe dispersion compensation techniques and the open problems facing them. Finally, Section 7 is speculation on the future of the field.

1. Dispersion in Optical Fibres

Fibres with core diameters of 50 and 62.5 μm are referred to as multimode fibre (MMF) because they can support many modes of light propagation at the optical wavelengths for which laser sources are readily available. Fibers with a core diameter of 10 μm support only one mode of propagation and are, hence, referred to as single-mode fibre (SMF). MMF is more mechanically robust and less expensive, but suffers from more dispersion than

SMF. Therefore, MMF is currently employed for links less than 1 km in length and data rates of 10 Gbps and lower, whereas SMF supports data rates of 40 Gbps over hundreds of kilometers. Currently, the rate and reach of both MMF and SMF links are limited by our ability to efficiently compensate for dispersion. The dominant forms of dispersion in commercial optical links today are chromatic dispersion and modal dispersion.

1.1 Chromatic Dispersion

Chromatic dispersion affects both MMF and SMF. It arises because different wavelengths of light propagate along optical fibers with different velocities. Any laser source modulated by random data produces light spread over a range of wavelengths. As a result, narrow light pulses broaden as they propagate along a fibre. Fortuately, chromatic dispersion can be inexpensively compensated for optically. For example, if the propagation delay of an optical fibre exhibits a known wavelength-dependency, it can be cascaded with a short section of fibre exhibiting the exact inverse wavelength-dependency. Ideally, the resulting link will be free from chromatic dispersion. Fibres can also be designed with propagation velocities that are nearly constant around the wavelength of a transmitting laser. Hence, chromatic dispersion is not the focus of most current research on integrated circuit dispersion compensation. Modal dispersion has proven more difficult to handle.

1.2 Modal Dispersion

Both MMF and SMF suffer from modal dispersion. The physical mechanisms are different, but both can result in similar time-varying channel responses.

In MMF links, a single transmitted light pulse excites multiple modes of propagation along the fibre, each experiencing a different delay and attenuation. Hence, each transmitted pulse results in multiple pulses of light at the receiver with different arrival times and amplitudes. This is similar to a wireless radio link with multipath fading, which is a notoriously difficult channel to communicate over. To complicate matters, energy from one propagation mode can spill over into another at discontinuities along the fibre [2]. Hence, a MMF channel response can change dramatically when a fibre or connector is mechanically stressed. Of course, there is a wealth of literature on how to handle multipath fading channels. Unfortunately, it is difficult to apply those same techniques to compensate for dispersion in MMF since the data rates in optical links are orders of magnitude greater than in our fastest wireless links.

Dispersion in MMF is reduced when the light source is not perfectly aligned in the center of the fibre core. In practice, this can be accomplished using a short section of fibre called a "patch cord" that provides an intentional offset between the light source and the center of the main fibre. This is a practical trick that mitigates, but does not completely eliminate, the problem. A more manageable pulse response results, often comprising only two pulses.

In SMF, although only a single mode of light propagation is supported, dispersion still arises due to asymmetries in the fibre cross-section. The asymmetries may be due to imperfections during manufacture, or due to mechanical stress after manufacture. Each pulse of light has energy in two orthogonal modes of polarization, which due to the asymmetries propagate with different velocities and attenuation. As a result, again, single optical pulses may split into two pulses by the time they reach the receiver. This phenomenon is referred to as polarization mode dispersion.

Dispersion due to multimode propagation in MMF and polarization mode dispersion in SMF are referred to, together, as *modal dispersion*. To get a feel for the seriousness of these channel impairments in practical applications, Figure 1 plots three typical pulse responses for a 10-Gbps link over 220 m of MMF with a patch cord.

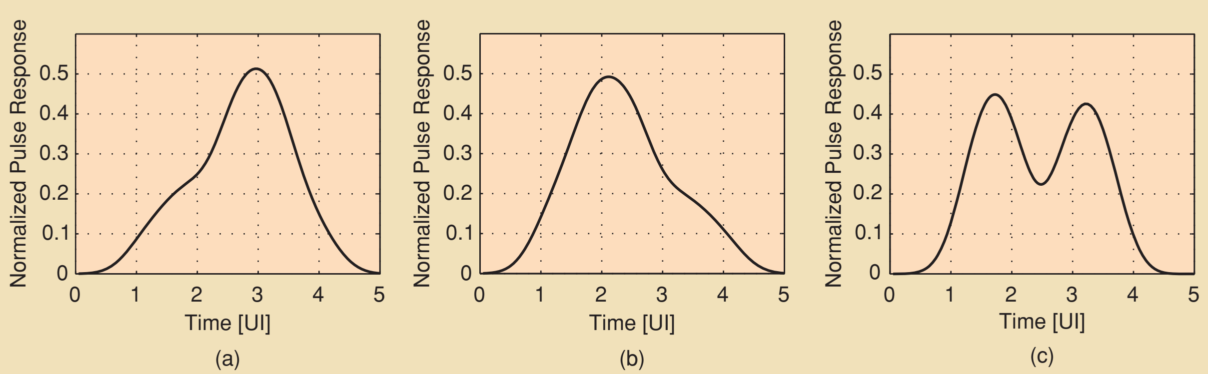

Figure 1. Example fibre channel pulse responses: (a) A precursor ISI pulse; (b) A postcursor ISI pulse; (c) A split pulse response.

In all three cases, the pulse response is confined to roughly 3 baud unit intervals (UI) or 300 ps at 10 Gbps. For the same length of fibre, SMF exhibits much less modal dispersion than MMF. Nevertheless, polarization mode dispersion has become a major issue in long-haul SMF links. Pulse responses similar to those in Figure 1 arise over hundreds of kilometers of SMF at 10 Gbps, and over tens of kilometers of SMF at 40 Gbps.

In Figure 1(a), fast mode of propagation is visible preceding the main pulse. Hence, in a binary communication stream each bit interferes with the two preceding bits. This is called a precursor response since it has mostly precursor intersymbol interference (ISI). The pulse response in Figure 1(b), on the other hand, has a slow mode of propagation following the main pulse resulting in mostly postcursor ISI.

Figure 1(c) is a particularly interesting case where the channel response has two distinct pulses that are nearly equal in amplitude and clearly separated in time. Any receiver for this channel will face some ambiguity in determining which pulse is the “main” one. We shall see that the resulting confusion necessitates nonlinear receiver architectures.

The normalized frequency responses corresponding to the three test cases in Figure 1 are plotted in Figure 2. All three have 3-dB bandwidths of around 2 GHz; not much for an “infinite-bandwidth”medium! Furthermore, modal dispersion can be vary significantly within a few milliseconds [3] demanding an adaptive approach to EDC.

The precursor and postcursor ISI responses in Figure 2(a) and 2(b) have a generally lowpass response. Hence, these types of dispersion can be mitigated by linearly filtering the received signal with the inverse channel response. Such linear equalization of the channel is considered in Section 2. However, the split pulse response in Figure 2(c) proves to be more problematic. Since the frequency response has deep nulls, a portion of the transmitted spectrum is essentially lost and can not be recovered by linearly filtering the received signal. Nonlinear methods for recovering the data in this case are discussed in Sections 3 and 4, digital implementations are described in Section 5, and diversity schemes are touched upon in Section 6.

2. When Linear Equalization is Sufficient, How Can it be Implemented with Low-Power?

A straightforward and potentially low-power method for mitigating ISI is to employ an adaptive linear filter that equalizes the channel response. Although it is possible to place the equalization filter at the transmitter, modal dispersion is time-varying and it is difficult to communicate information about the channel from the receiver back to a transmit equalizer in real time. Hence, equalization for EDC is generally performed at the receiver.

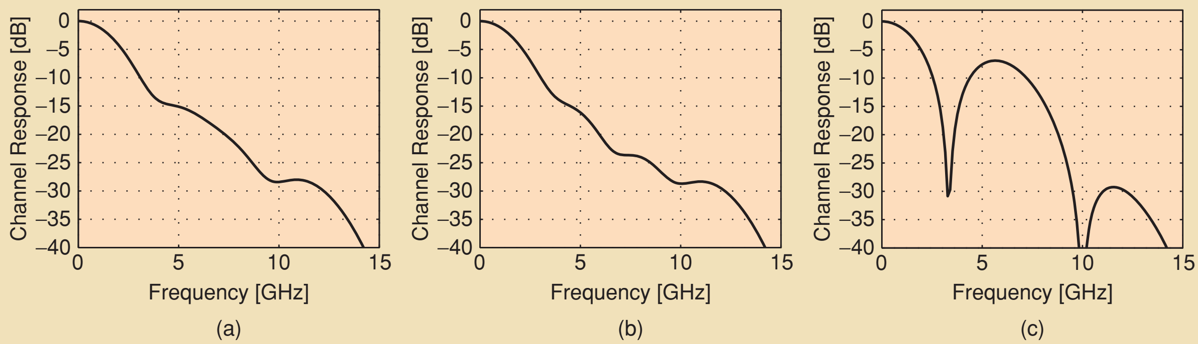

Figure 2. Example fibre channel frequency responses: (a) A precursor ISI pulse; (b) A postcursor ISI pulse; (c) A split pulse response.

Simple analog equalizers with programmable high-frequency boost such as [4] have been implemented with relatively low power at data rates up to 10 Gbps. Although these work well for copper links, they do not provide sufficient flexibility for EDC. We have already seen that modal dispersion can result in a wide variety of channel responses sometimes having deep nulls at unpredictable frequencies.

Programmable finite impulse response (FIR) filters can accommodate the wide variety of fibre responses attributable to modal dispersion, making them a popular choice for EDC applications. FIR equalizers have long been employed in magnetic storage applications at data rates exceeding 1 Gbps. The filters are guaranteed stable, and established techniques exist for the adaptation of their coefficients. However, their low-power implementation remains an open research topic.

An FIR filter requires analog delays. A clock can be used to define the delay intervals [5], but its distribution with low-skew and low-jitter generally burns a lot of power. Alternately, analog continuous-time delays may be synthesized, for example using LC-ladders with integrated inductors as delays. For this purpose, integrated circuit designers have recently rediscovered a filter topology that was originally demonstrated over 35 years ago in a discrete prototype [6]. The traveling wave filter topology shown in Figure 3 cleverly distributes the tap amplifiers' parasitic capacitances along two LC delay lines, similar to a distributed amplifier. However, unlike a distributed amplifier, the output is taken at the near end of the transmission line at the top of Figure 3 so that the tap amplifier outputs are summed out-of-phase. Hence, the broadband delay introduced by the passive LC networks can be used to create a finite impulse response that is programmable via the amplifier gains, p_k. Unfortunately, the series resistance, skin effect, and substrate losses associated with integrated spiral inductors lead to signal attenuation that increases with the length of the line and with frequency. This has prompted more than one frustrated researcher to note that such equalizers need equalizers! Instead, active circuits can be used to provide a continuous-time delay without loss [7] but they also tend to consume a lot of power, especially when designed for high bandwidths.

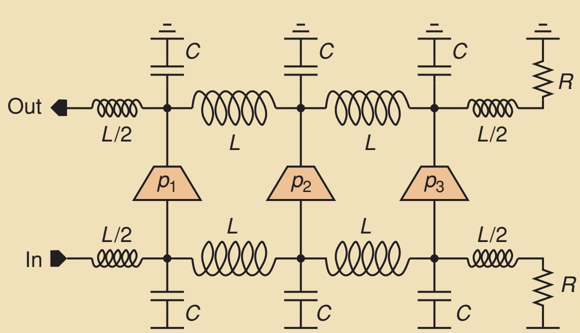

Figure 3. A 3-tap traveling wave filter.

Research on broadband FIR equalizers is ongoing and active. Novel topologies that distribute parasitic capacitances along LC delay lines have been developed to provide higher bandwidth [8], [9] and low power [10]. Others have employed small active circuits to mitigate losses in passive delay lines [11]. Research on these topics will no doubt continue, but for longer links and/or higher data-rates, more sophisticated methods for EDC are being researched.

Ultimately, if the channel has deep nulls in its frequency response, linear equalization is insufficient. Some form of nonlinear processing is required to restore the lost portion of transmitted spectrum. The most common examples are decision feedback equalization, and maximum likelihood sequence estimation.

3. How Can the Timing Requirements of Decision Feedback be Met?

A decision feedback equalizer (DFE) is a nonlinear receiver that can restore lost portions of transmit spectrum using past decisions. A DFE combines a feedforward linear equalizer and a feedback equalizer whose input is the output of a nonlinear decision circuit.

DFEs are very flexible and capable of handling any pulse response so long as the number of taps is sufficient. However, their speed is limited by a critical timing path through the feedback filter. This path must settle within a bit period: only 25 ps for a 40-Gbps link.

To shorten this critical path and enable higher speed operation, parallel processing can be employed in a "speculative" or "lookahead" DFE architecture [12] with increased circuit power and area. Such architectures are practical for one or two taps and have been demonstrated up to 40 Gbps [13].

Parallelism is, of course, costly—it increases the power consumption and area of the integrated circuits, thus increasing cost. Parallel architectures also generally complicate adaptation of the EDC and timing recovery. Nevertheless, DFEs are currently the focus of most commercial EDC development. Meanwhile, further performance improvements are already being sought.

4. MLSE is a Clear Winner on Paper, But Implementation Above 10 Gbps Remains an Open Problem

Since a DFE cancels post-cursor ISI, it essentially "throws away" signal energy. A better receiver will recognize specific patterns of ISI and use them to inform its decision. Maximum likelihood sequence estimation (MLSE) does exactly this and thereby provides better performance than a DFE in dispersive optical links. For example, Figure 4 shows a performance improvement of roughly 1.5-dB for an ideal implementation of MLSE compared with a 2-tap DFE. However, the complexity of MLSE is significant. The simulated results assume perfect a *priori* knowledge of the channel response at the receiver. They also assume that energy from the entire pulse response was used for detection which would imply massive complexity in implementing MLSE.

A partial response maximum likelihood (PRML) receiver employs a linear equalizer preceding MLSE to partially equalize the channel response. The linear equalizer shortens the channel's pulse response to only a few UI in duration. The following MLSE now works on the shortened partial response channel with practical complexity. Since the linear equalizer introduces some noise amplification and correlation, the performance of a PRML receiver lies somewhere between that of the DFE and ideal MLSE in Figure 4.

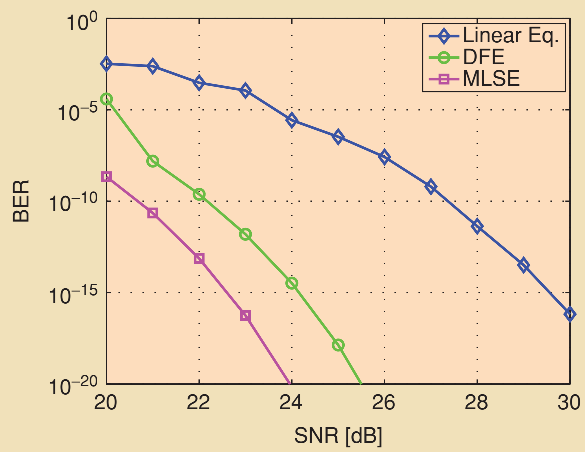

Figure 4. Bit error rate versus SNR for the split-pulse channel response in Figure 1(c) using a 6-tap 1/2-UI-spaced linear equalizer, a DFE comprising the same 6-tap 1/2-UI-spaced linear equalizer with a 2-tap feedback equalizer, and a MLSE receiver.

However, if the partial response target can be chosen to closely match the channel's inherent response, the noise amplification is minimized and performance close to that of ideal MLSE is obtainable.

The Viterbi algorithm is a practical realization of MLSE. Analog implementations of the Viterbi algorithm have been reported for relatively short partial response targets at data rates of hundreds of Mbps [14]. The implementation of an analog MLSE receiver in CMOS suitable for the long and varied channel responses encountered in, for example, 200 m MMF links at 10 Gbps is a noble and lofty research goal indeed. Sadly (for die-hard believers in analog signal processing like me), most research effort on MLSE receivers at 10 Gbps has focused on digital implementations.

5. When Will Advances in Data Converters Open the Door to DSP-Based Transceivers?

Integrated circuit transceivers for most every digital communication application in the past 20 years have matured, predictably, from mostly-analog architectures towards mostly-digital ones, and there is little reason to think that optical links with EDC will be any different. Voice-band modem receivers once employed analog bandpass filters to detect frequency-modulated signals, but eventually comprised simply an analog-to-digital converter (ADC) followed by sophisticated digital signal processing (DSP) described in software. A similar story is currently being played out in many wireless communication applications, where DSP is annexing more and more transceiver functionality as quickly as advances in CMOS ADCs permit.

Magnetic storage read channels are a particularly interesting case-in-point. Analog Viterbi detection was the subject of intense research effort in the late 1990's. But the industry appears to have completely lost interest in it, notwithstanding the protestations of countless analog signal processing researchers. Two factors contributed to this trend. First, analog Viterbi detectors were not well suited to the longer and adaptive partial response targets sought to increase storage density. Second, advancing CMOS process technologies favor mostly-digital architectures. So nowadays a typical read channel architecture features a front-end ADC followed by an impressive custom DSP. Industry's brief interest in analog Viterbi detection did gave a boost to research on analog error control coding, but to this day analog decoding remains a solution in search of a problem.

Hence, it would seem that a digital receiver is the ultimate solution for EDC as well. If a suitable ADC can be provided at the receiver front end, any of the preceding approaches to EDC can be applied. Digital MLSE receivers have already been reported [15]. The current state-of-the-art provides 4 to 6 bit ADCs for 10 Gbps systems, but research towards 40 Gbps systems is ongoing, with front-ends for the required ADC having already been demonstrated [16].

Placing an ADC at the front-end of an optical receiver opens the door to a host of exciting research opportunities to exploit sophisticated signal processing for further increases in rate, reach, and robustness. For example, OFDM is an effective alternative to equalization for mitigating multipath ISI. OFDM is already popular for wireless channels at lower data rates, and has been studied for optical links up to 40 Gbps [17]. The "soft" decisions provided by a front-end ADC can also be used to improve the performance of error control codes. Low density parity check codes with soft-decision decoding have been suggested for optical links [18] with performance approaching the theoretical limit of channel capacity.

6. What is the Future Role of Adaptive Optics in Dispersion Compensation?

Extending the analogy between modal dispersion in optical fibers and multipath fading in wireless receivers, one might be inspired to consider diversity schemes to combat modal dispersion. It is well known that multiple transmit and receive antennae can improve wireless reception in a fading environment. Similarly, multiple lasers and photodetectors can be employed to turn a single optical fibre into a multiple-input multiple-output (MIMO) communication channel with an attendant increase in channel capacity [19].

At the transmitter, diversity can be used to spatially modulate the light source [20]. If done correctly, the launch energy can be completely confined to only one mode of propagation even in a MMF thereby providing greatly improved signal integrity. At the receiver, the outputs of an array of photodetectors can be combined to cancel interfering propagation modes for a kind of "spatial equalization" [21]. Whether at the transmitter or receiver, the optics must be adapted to track time-variations in the channel.

Although promising, this avenue of research is still in early stages. Segmented light sources and detectors that are robust and low-cost are still needed. Advances in signal processing are also required since MIMO methods are currently limited to wireless communication operating at a small fraction of the data rates in optical fibre links.

7. The Light at the End of the Tunnel

This article just scratches the surface on EDC, and several difficult challenges are not even touched upon. For instance, how can EDC be made adaptive in real time? How can EDC be efficiently combined with timing recovery? But even this cursory survey shows that light has its limits and current research is exploring ever-more sophisticated (i.e. *expensive*) integrated circuit signal processing to push those limits. In light of all these challenges (no pun intended) it is tempting to throw up our arms, give up and conclude that optical links are doomed by dispersion.

But optical links have some excellent properties that will not let us set them aside. For example, unlike electrical links, optical links have very low loss. Most of the transmitted energy eventually arrives at the receiver, albeit a little garbled by dispersion. Furthermore, relatively little noise is introduced along fibre, and fibers can be bundled tightly together with much less crosstalk than copper wires. So if dispersion can be managed, the potential for very broadband and long reach communication exists. Perhaps solutions lie in clever combinations of optical and electronic signal processing. One thing is certain: overcoming the limits of light cost-effectively will require a circuits AND systems approach.

References

[1] R. Palmer, J. Poulton, T. Greer, M. Kellam, F. Quan, F. Zarheshvari, and A. Fuller, "A 14mW 6.25Gb/s transceiver in 90nm CMOS for serial chip-to-chip communication," in *International Solid-State Circuits Conference, Digest of Technical Papers,* San Francisco, CA, 2007, pp. 440–441.

[2] P. Pepeljugoski, S.E. Golowich, A.J. Ritger, P. Kolesar, and A. Risteski, "Modeling and simulation of next-generation multimode fiber links," *J. Lightwave Technol.,* vol. 21, no. 5, pp. 1242–1255, May 2003.

[3] H. Bülow, W. Baumert, H. Schmuck, F. Mohr, T. Schulz, F. Kuppers, and W. Weiershausen, "Measurement of the maximum speed of PMD fluctuation in installed field fiber," in *Optical Fiber Communication Conference, 1999, and the International Conference on Integrated Optics and Optical Fiber Communication. OFC/IOOC '99. Technical Digest,* vol. 2, San Diego, CA, Feb. 1999, pp. 83–85.

[4] Y. Tomita, M. Kibune, J. Ogawa, W. Walder, H. Tamura, and T. Kuroda, "A 10-Gb/s receiver with series equalizer and on-chip ISI monito in 0.11-μm CMOS," *IEEE J. Solid-State Circuits,* vol. 40, no. 4, pp. 986–993, Apr. 2005.

[5] J. Jaussi, G. Balamurugan, D. Johnson, B. Casper, A. Martin, J. Kennedy, N. Shanbhag, and R. Mooney, "8-Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation, and clock de-skew," *IEEE J. Solid-State Circuits,* vol. 40, no. 1, pp. 80–88, Jan. 2005.

[6] W. Jutzi, "Microwave bandwidth active transversal filter concept with MESFETs," *IEEE Trans. Microwave Theory Tech.,* vol. MTT-19, no. 9, pp. 760–767, Sept. 1971.

[7] D. Harnandez-Garduno and J. Silva-Martinez, "A CMOS 1Gb/s 5-tap transversal equalizer based on inductorless 3rd-order delay cells," in *International Solid-State Circuits Conference, Digest of Technical Papers,* San Francisco, CA, Feb. 2007, pp. 232–233.

[8] J. Sewter and A. Chan Carusone, "A CMOS finite impulse response filter with a crossover traveling wave topology for equalization up to 30 Gb/s," *IEEE J. Solid-State Circuits,* vol. 41, no. 4, pp. 909–917, Apr. 2006.

[9] J. Sewter and A. Chan Carusone, "A 3-tap FIR filter with cascaded distributed tap amplifiers for equalization up to 40 Gb/s in 0.18-μm CMOS," *IEEE J. Solid-State Circuits,* vol. 41, no. 8, pp. 1919–1929, Aug. 2006.

[10] R. Tiruvuru and S. Pavan, "Transmission line basd FIR structures for high speed adaptive equalization," in *IEEE Int. Symp. Circuits and Systems,* Kos, Greece, May 2006, pp. 1051–1054.

[11] S. Reynolds, P. Pepeljugoski, J. Schaub, J. Tiemo, and D. Beisser, "A 7-tap transverse analog-fir filter in 0.13 μm CMOS for equalization of 10Gb/s fiber-optic data systems," in *IEEE ISSCC Dig. Tech. Papers,* San Francisco, CA, Feb. 2005, pp. 330–331.

[12] S. Kasturia and J.H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Selected Areas in Communications*, vol. 9, no. 5, pp. 711–717, June 1991.

[13] A. Garg, A. Chan Carusone, and S.P. Voinigescu, "A 1-tap 40-Gbps look-ahead decision feedback equalizer in 0.18-μm SiGe BiCMOS technology," *IEEE J. Solid-State Circuits,* vol. 41, no. 10, pp. 2224–2232, Oct. 2006.

[14] B. Zand and D. Johns, "High-speed CMOS analog Viterbi detector for 4-PAM partial-response signaling," *IEEE J. Solid-State Circuits,* vol. 37, no. 7, pp. 895–903, July 2002.

[15] H. min Bae, J. Ashbrook, J. Park, N. Shanbhag, A. Singer, and S. Chopra, "An MLSE receiver for electronic-dispersion compensation of OC-192 fiber links," in *International Solid-State Circuits Conference, Digest of Technical Papers,* San Francisco, CA, Feb. 2006, pp. 874–883.

[16] S. Shahramian and A. Chan Carusone, "Design methodology for a 40-GSamples/Sec track and hold amplifier in 0.18-μm SiGe BiCMOS technology," *IEEE J. Solid-State Circuits,* vol. 41, no. 10, pp. 2233–2240, Oct. 2006.

[17] K. Azadet and F. Saibi, "Bandwidth efficient modulation over fiber optics," in *Compound Semiconductor Integrated Circuits Conference,* Palm Springs, CA, Oct. 2005, pp. 185–188.

[18] B. Vasic, I. Djordjevic, and R. Kostuk, "Low-density parity check codes and iterative decoding for long-haul optical communication systems," *J. Lightwave Technol.,* vol. 21, no. 2, pp. 438–446, Feb. 2003.

[19] H. Stuart, "Dispersive multiplexing in multimode optical fiber," *Science,* vol. 289, no. 5477, pp.281–283, July 2000.

[20] J. Kahn, "Compensating multimode fiber dispersion using adaptive optics," in *Optical Fiber Conference,* Anaheim, California, Mar. 2007, p. OTuL1.

[21] K. Patel, A. Polley, K. Balemarthy, and S. Ralph, "Spatially resolved detection and equalization of modal dispersion limited multimode fiber links," *J. Lightwave Technol.,* vol. 24, no. 7, pp. 2629–2636, July 2006.

Anthony Chan Carusone completed the B.A.Sc. and Ph.D. degrees at the University of Toronto in 1997 and 2002 respectively. Since 2001, he has been a faculty member in the Department of Electrical and Computer Engineering at the University of Toronto where he is currently an Associate Professor. Dr. Chan Carusone was a co-author of the best paper at the 2005 Compound Semiconductor Integrated Circuits Symposium and the best student paper at the 2007 Custom Integrated Circuits Conference. He is a past chair of the Analog Signal Processing Technical Committee for the IEEE Circuits and Systems Society, a member of the technical program committee for the Custom Integrated Circuits Conference, and Deputy Editor-in-Chief for the *IEEE Transactions on Circuits and Systems II: Express Briefs*.