Lighting the Local Area: Optical Code-Division Multiple Access and Quality of Service Provisioning

Andrew Stok and Edward H. Sargent, University of Toronto

Abstract

Optical code-division multiple-access is proposed as a natural solution to achieving asynchronous, high-speed connectivity in a local area network environment. Optical CDMA is shown to be competitive with other networking technologies such as WDMA and TDMA, but has the benefit of more flexibility, simpler protocols, and no need for centralized network control. The limitations of one-dimensional optical orthogonal codes for CDMA have motivated the idea of spectral spreading in both the temporal and wavelength domains. If the constraints on constant weight in these two-dimensional codes are relaxed, differentiated levels of service at the physical layer become possible. Areas for further research are suggested which may allow quality of service levels to be guaranteed at the physical layer.

The success of long-haul fiber optic communication systems has shifted the focus of optical networking to the shorter-haul metropolitan and local-area domains. Now that vast quantities of information can be sent over thousands of kilometers at low cost, the limiting path will be in the local area and accessing the long distance light pipe. While the performance of long-distance dense wavelength-division multiplexed (DWDM) systems is announced in terabits per second, most of us continue to use networks whose aggregate capacities are 10 or 100 Mb/s. Standards and systems have recently been announced to achieve networks with up to 10 Gb/s aggregate capacities, but this may not be enough. A local area network (LAN) may well comprise hundreds of users, each of whom may employ data visualization, high-definition digital video broadcasts, or other bandwidth-intensive applications. Each user may well require individual data rates in gigabits per second, leading to aggregate data rates reaching hundreds of gigabits per second.

It will not be sufficient to provide raw bandwidth alone. The network must also provide quality of service (QoS) guarantees for these applications, even as the number of users and aggregate throughput change with time. This challenge is not insignificant — the network must simultaneously accommodate traffic whose requirements vary over orders of magnitude:

- Bandwidths ranging from kilobits (compressed voice) to gigabits per second (high-quality, dynamic, three-dimensional, real-time image sharing)
- Bit error rates ranging from 10^{-3} (Internet telephony) to 10^{-12} (highly sensitive data transfer)
- Delays/latencies ranging from speed-of-light-limited shared memory applications (e.g., less than 1 μs in a LAN of modest size) to file transfer (e.g., seconds to minutes)

There is a worrisome conflict between the required speed for data transmission and networking to occur, and the sophistication of differentiated service provisioning demanded.

Sharing the Fiber Optic Medium Using Codes

We need to find a way to share the vast bandwidth of the optical fiber medium in a manner which is fair and fast. Three multiple access approaches are often considered to make the system bandwidth available to the individual user: time-division multiple access (TDMA), wavelength-division multiple access (WDMA), and code-division multiple access (CDMA).

TDMA accommodates a large number of active users by interleaving bits from different sources into a period equal to that of the uncompressed bit. TDMA systems offer a large number of node addresses; however, the performance of TDMA systems is ultimately limited by the time-serial nature of the technology. The receiver must operate at the aggregate bit rate of the system, which is roughly equal to the number of nodes connected times the data rate per node. Even a modest system with only 100 users each communicating at 1 Gb/s would result in a total throughput well beyond the typical electronic hardware speeds of several gigabits per second.

TDMA systems also require strong centralized control to allocate time slots and maintain synchronous operation. Although centralized control is straightforward to implement

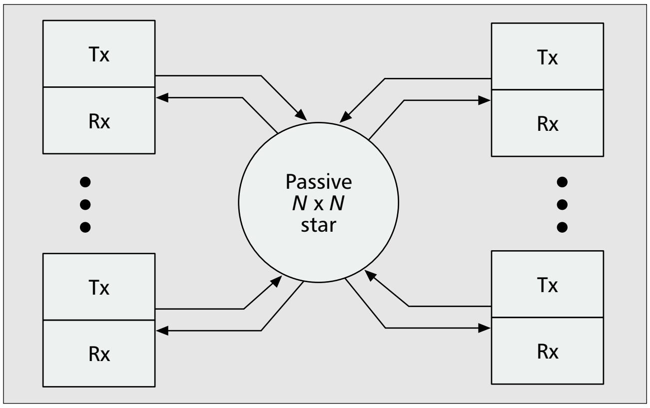

■ Figure 1. *A fiber optic CDMA network using a passive* N x N *star coupler.*

in a LAN, the extra time required for users to request time slots and for the central node to allocate the slot may introduce significant latency. Furthermore, the allocation of dedicated time slots does not allow TDMA to take full advantage of statistical multiplexing gain, which may be significant when the data traffic is bursty.

In the WDMA approach, the available optical bandwidth is divided into distinct wavelength channels that are used concurrently by different users to achieve multiple access. If each wavelength channel were to support a data rate equal to the peak electronic capacity of several gigabits per second, a WDMA system with around 160 wavelength channels could accommodate on the order of a terabit per second of aggregate capacity.

The problem with using WDMA in LANs is that a significant amount of dynamic coordination between nodes is required. A dedicated control channel can be used for pretransmission coordination; however, this wastes bandwidth that could otherwise be used for data transmission and introduces latency as nodes attempt to negotiate a connection. The control channel can be avoided by assigning transmission rights in a predetermined fashion, or through contention-based protocols [1]. However, in large networks with dynamic populations scheduling becomes difficult, and collision detection schemes waste bandwidth and can introduce significant latency.

Optical CDMA offers an interesting alternative for LANs because neither time management nor frequency management of all nodes is necessary. Optical CDMA can operate asynchronously, without centralized control, and does not suffer from packet collisions; therefore, very low latencies can be achieved. Dedicated time or wavelength slots do not have to be allocated, so statistical multiplexing gains can be high. In contrast to TDMA and WDMA where the maximum transmission capacity is determined by the total number of these slots (i.e., hard-limited), CDMA allows flexible network design because the bit error rate (BER) depends on the number of active users (i.e., soft-limited).

Optical CDMA has the additional advantage of providing differentiated service at the physical layer, unlike TDMA or WDMA where QoS is usually handled in the higher open systems integration (OSI) levels such as the data link or network layer. A user needing more bandwidth can be allocated more time slots in TDMA or wavelength channels in WDMA; however, these schemes may waste bandwidth when the traffic is bursty, and there is no mechanism to allow different users to transmit with different BERs. In this article an extension of optical CDMA called multiple-wavelength optical CDMA will be shown to provide differentiated service in the physical layer. Implementing some QoS in lower OSI layers may reduce overhead and simplify LAN design.

Optical Code-Division Multiple Access

Optical CDMA combines the large bandwidth of the optical medium with the flexibility of CDMA to achieve high-speed connectivity. CDMA was originally investigated in the context of radio frequency (RF) communications systems, and was first applied to the optical domain in the mid-1980s by Prucnal, Salehi, and others [2–4]. These researchers sought to use the excess bandwidth in single-mode fibers to map low-information-rate electrical signals into high-rate optical pulse sequences to achieve random asynchronous operation without the need for a centralized network controller. A typical fiber optic CDMA system is shown in Fig. 1, where the nodes are connected together through a passive $N \times N$ star coupler. At the logical level this configuration is a broadcast-and-select network.

In an optical CDMA system, each bit is divided up into n time periods, called *chips*. By sending a short optical pulse during some chip intervals, but not others, an optical signature sequence, or *codeword*, can be created. The total number of illuminated chips in the address codeword is called the (Hamming) weight w. Each user on the optical CDMA system has a unique signature sequence. The encoder of each transmitter represents each 1 bit by sending the signature sequence; however, a binary 0 bit is not encoded and is represented using an all-zero sequence. The encoded signal is sent to the $N \times N$ star coupler and broadcast to all nodes. The crosstalk between different users sharing a common fiber channel, known as *multiple-access interference* (MAI), is usually the dominant source of noise in an optical CDMA system; therefore, intelligent design of the codeword sequences is important to reduce the contribution of MAI to the total received signal.

Optical Orthogonal Codes

The key to an effective CDMA system is the choice of efficient address codes with good correlation properties for encoding the source bits. The problem of optimal code selection has been addressed in the context of wireless microwave communications in which CDMA is one of the dominant technologies. In RF CDMA systems, sequences consist of $\{-1, +1\}$ values and are therefore known as *bipolar codes*. The choice of both positive and negative bit values is natural since phase of the electromagnetic field can be detected directly. However, in optical systems the signal energy, rather than amplitude and phase, is most readily measured at the detector. In this case, the direct application of bipolar codes to optical CDMA communications is not possible due to the nonnegative nature of incoherent photodetection. As a result, only *unipolar* sequences consisting of $\{0, 1\}$ values can be used for optical CDMA systems. In response, new families of CDMA codes called *optical orthogonal codes* (OOCs) have been designed specifically for unipolar systems [3]. Two OOC address sequences of length $n = 32$ and weight $w = 4$ are shown in Fig. 2.

One of the primary goals of optical CDMA is to extract data with the desired address code in the presence of all other users' optical pulse sequences; therefore, the set of codewords should be designed to satisfy three fundamental conditions:

- For any codeword the nonshifted autocorrelation, equal to the Hamming weight of the codeword, should be made as large as possible. This ensures that the received signal is much larger than the background noise in the system.
- For any codeword the shifted autocorrelation must be much less than the Hamming weight. This requirement ensures that the output of the correlator receiver will be small when the receiver is not synchronized with the transmitter, and allows optical CDMA to operate asynchronously without the need for a global clock signal.

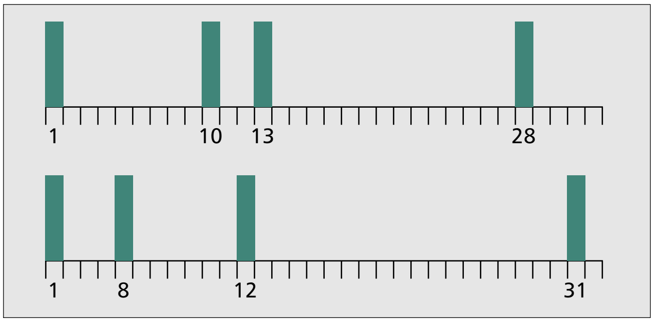

■ Figure 2. *Two optical orthogonal codes with length* n = *32, weight* w = *4, and autocorrelation and cross-correlation constrained to one.*

- The cross-correlation between any pair of codewords must small. This property ensures that each codeword can easily be distinguished from every other address sequence. In other words, we seek to make the MAI insignificant compared to the energy contained in the received information bit.

A set of address codes that satisfy the three correlation conditions listed above will allow asynchronous operation of the system and minimize the BER by managing the MAI noise term.

There is an important additional consideration in any shared medium network. Since only a fraction of the nodes on a network will be transmitting at any given time, it is necessary to accommodate many more total network members than simultaneous network users. Since each network member must be assigned a distinct code with all of the desired correlation properties, the size of the set of all allowed codewords — the *cardinality* of the code — must meet or exceed the number of network members to be supported. The cardinality of a unipolar code set which meets specified correlation properties is much less than the cardinality of a corresponding set of bipolar codes. Bipolar codes can be constructed such that the signals add to zero, leading to lower cross-correlation values and hence a larger number of available address codes; this cancellation is not possible in positive systems.

The upper bound on the cardinality Φ of a set of CDMA codes with unity autocorrelation and unity cross-correlation is [3]

$$\Phi(n,w) \le \left\lfloor \frac{n-1}{w(w-1)} \right\rfloor \quad (1)$$

where the symbol $\lfloor z \rfloor$ denotes the integer portion of the real value z, n is the code length, and w is the code weight. To maximize the number of users that can be served simultaneously, the length should be made large and the weight small. Effective optical orthogonal codes are thus very sparse. This has two important consequences for CDMA system design. First, the energy per encoded bit is low and may compromise the overall energy budget of the system. Second, very sparse codes imply that the chip period must be much smaller than the source bit length. Very high speed (and hence expensive) electronic and optical equipment will be required to produce these very short pulses. Furthermore, as the duration of each chip becomes shorter, dispersion effects could start to limit performance severely.

The code weight also has a direct effect on the performance of an optical CDMA system. If the code weight were increased but the threshold level kept low, system performance would degrade since, by increasing the number of pulses per frame, one increases the probability of multiple codewords overlapping in the same chip. To improve system performance, both the code weight and threshold should be increased. It is then less probable that multiple users will occupy the same chip up to the level of the threshold [4], so MAI is reduced.

Higher-Dimensional Optical Codes: Exploiting Additional Degrees of Freedom

For the optical orthogonal codes discussed in the previous section, good autocorrelation and cross-correlation properties were achieved through the use of very long signature sequences. The resulting bandwidth expansion leads to stringent requirements on the speed of the encoding and correlating hardware.

The technological demands can be reduced by taking advantage of an extra degree of freedom: the wavelength of light. In 1988 Perrier and Prucnal proposed combining multiple wavelength channels with optical CDMA in a scheme called WDMA + CDMA [5]. Every signature sequence is reusable and is sent simultaneously at different wavelengths. The code length and hence the hardware speed are reduced, since only a small fraction of the original number of users share a single wavelength channel. The WDMA + CDMA scheme required a centralized network controller to distribute wavelength channels uniformly to the transmitting nodes. It was therefore difficult to allocate the possible wavelengths efficiently, and the scheme failed to achieve the expected performance [6].

An alternative approach to reducing the demands on the electronic hardware is to spread the optical orthogonal codes in both the temporal and wavelength domains simultaneously. Instead of viewing each wavelength as a separate channel that can support a set of optical orthogonal codes, the time chips and wavelength channels can be viewed as the axes of a two-dimensional codeword. Figure 3 demonstrates how the multiple-wavelength optical CDMA scheme compares to conventional WDMA and TDMA approaches.

Analogous to the one-dimensional OOCs, these multiple-wavelength OOCs (MWOOCs) are designed to meet autocorrelation and cross-correlation constraints in order to manage MAI. MWOOCs can in general support a much larger number of simultaneous users since the code cardinality is much larger. For example, an optical CDMA system using a code set with length $n = 1000$ and code weight $w = 9$ has an upper bound of 13 simultaneous users, according to Eq. 1. Using an upper bound on cardinality for two-dimensional codes derived by Yang and Kwong [6], it is easy to see the difference an extra degree of freedom makes. Adding just three more wavelength channels to the system would allow a set of 1000 x 4 MWOOC codewords to be used, and the upper limit on the number of simultaneous users would increase to over 1300. This increase in the maximum network size does not require a corresponding increase in the electronic processing speed of the system since both the 1D and 2D codewords have the same temporal length of $n = 1000$.

In moving from single-wavelength optical CDMA systems to multiple-wavelength networks, there is the danger of introducing additional latency. To keep costs low, it may be desirable to use a single tunable laser to produce the chips at multiple wavelengths; however, at high data rates, tuning the laser from one chip to the next at a different wavelength may not be possible within a chip period. A guard time could be added to allow the laser to settle to the desired wavelength before the pulse is required; however, this would reduce overall network speed and complicate operation since the settling time would depend on the "distance" in wavelength between two adjacent chips.

The guard time may be eliminated if a bank of fixed-wavelength lasers, each tuned to one of the required wavelengths, were provided in each transmitter. The performance of the network is improved, but the cost of the system is increased. Alternatively, a single broadband source may be spectrally sliced to provide all the required wavelength channels. With a large wavelength spacing (i.e., 200 GHz) lower-tolerance spectral filters could be used to keep costs low.

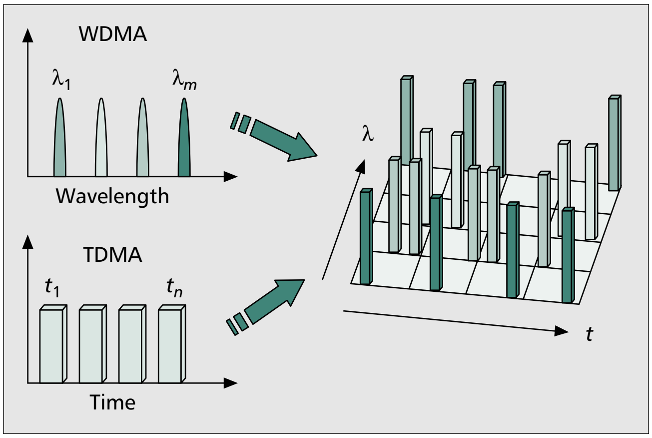

■ Figure 3. *A graphical representation of WDMA, TDMA, and multiple-wavelength optical CDMA. TDMA and WDMA only take advantage of a single degree of freedom (time and wavelength, respectively). The multiple-wavelength optical CDMA scheme spreads data in both the time and wavelength domains to achieve greater flexibility and eliminate centralized network control.*

The MWOOCs in general perform better than WDMA + CDMA schemes for a given bandwidth expansion and number of wavelengths. Designing 2D signature sequences, rather than artificially partitioning the available bandwidth into independent wavelength channels, leads to a lower probability of error. For example, Yang has shown that moving to MWOOCs from the 1D codes allowed for a large number of simultaneous users (approximately 600) with a very low probability of error (10^{-8}) for a reasonable codeword size of 31 x 31 [6]. Tancevski and Andonovic investigated 2D codewords that used extended quadratic congruence/prime sequences to spread bits in both the time and wavelength domains. A 31 x 31 codeword could support a system with 600 users and achieve a BER of approximately 10^{-7} [7].

Various techniques have been developed to construct MWOOCs such as modification of 1D OOCs, application of error correction codes to the optical domain, or algebraic coding methods. In all cases, however, the codes were arbitrarily restricted to have a single pulse per row in the matrix. This limited the code cardinality and necessitated equal weight for all signature sequences. As we show in the following section, relaxing the single pulse per row constraint will lead to a more flexible family of MWOOCs that may allow for differentiated levels of service at the physical layer in a CDMA LAN.

Multiple-Weight Two-Dimensional Codes

In the late 1990s, Kwong and Yang developed a new class of MWOOCs without the restrictions of one constant code weight [8, 9]. Through the algebraic code generation technique, two different classes of codewords were created: low weight w and high weight $2w$. These multiple-weight MWOOCs are the key to providing different QoS levels in an optical CDMA network *at the physical layer* in the OSI model. Since the performance of a signature pattern varies with its weight, multiple code weights enable a single network to support multiple performance requirements. A user with demanding service requirements can be assigned a signature pattern with a heavier weight in order to guarantee a lower error probability. Lower-priority users are then assigned the lighterweight codes.

Figure 4 shows the error probability as a function of block length for a CDMA system with multiple-weight 2D codewords. The number of users in the system is represented by $K = K_0 + K_1$ where K_0 users have weight w_0 codes and K_1 users have weight w_1 codes. The code weight of the multiple-weight MWOOC is a more important factor determining the performance, while the number of simultaneous users plays a lesser role. For example, the two middle curves in Fig. 4 show that users with codewords of weight 6 always perform better than those with codewords of weight 3, even though the former has more simultaneous users [6].

From Differentiated Service to QoS

The MWOOCs discussed in the previous section allow for differentiated levels of service at the physical layer but do not allow QoS guarantees to be provided. A critical user with a higher code weight will always have better performance than lower-weight users; however, no node on the LAN can be assured that the level of service will not fluctuate as network conditions change. The network must be able to compute not only the guaranteed end-to-end QoS for the newly admitted session, but must also compute the impact of this admission on already accepted sessions in the network.

The solution may be to design a CDMA network with dynamically allocated signature sequences to compensate for degraded QoS when new users are admitted. Should the network detect high MAI levels that compromise previous QoS guarantees, new signature sequences could be allocated to all nodes to satisfy both new and existing QoS contracts.

This scheme might appear to require strong centralized network control to allocate new codes to users and monitor the total MAI noise on the channel; however, by suitably modifying the CDMA encoders, the control can be distributed throughout the network. As shown in Fig. 5, each node could be equipped with a simple and inexpensive photodiode that measures the total energy on the channel, plus some code generation logic. Should the average MAI power increase beyond a predetermined threshold, each node would generate a new encoder signature sequence to maintain the desired QoS parameters. Each node would be assigned a distinct signature sequence for each of the possible code sets that may be encountered to avoid having two nodes transmit with the same address code.

Distributed network management would allow QoS to be provided at the physical layer, without any centralized control. An advanced code generation algorithm is needed, one that differs from the common algebraic coding techniques. The code generation algorithms available today (e.g., the MWOOCs of Kwong

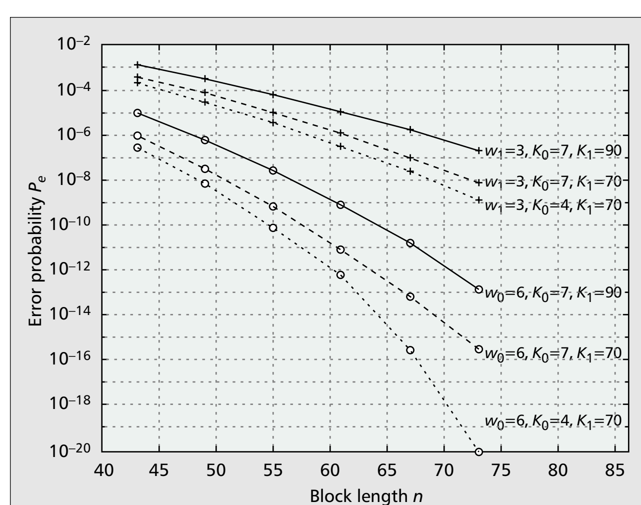

■ Figure 4. *Error probability as a function of block length* n *for the multiple-weight MWOOC. The two code weights are* $w_0 = 6$ *and* $w_1 = 3$ *[8].*

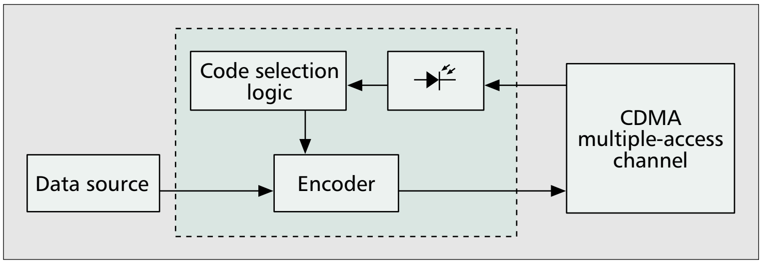

■ Figure 5. *A dynamic encoder that adjusts the signature sequence based on the total network traffic.*

and Yang [8]) produce a set of signature sequences with various weights. Numerical simulations are then used to determine the BERs that can be supported by the network with these codes. However, for the distributed network management scheme, the required BER and number of users must be the input to the code generation algorithm, and the output should be a set of signature sequences that can support the desired QoS levels. Simulation should not be required to determine which QoS levels a signature sequence can support, given the background noise. Unfortunately, the QoS that can be supported by a code is not a linear function of the code weight, which complicates the algorithm. The design of an advanced code generation algorithm is a possible direction for future research to allow provision of QoS in optical CDMA LANs at the physical layer.

Opportunities, Directions, and Conclusions

Optical code-division multiple access has come a long way since the seminal research of the 1980s. Nevertheless, a host of opportunities remain for achieving truly novel network functionality at the optical layer.

The combination of optical CDMA with forward error correction may dramatically increase the efficiency with which the vast optical bandwidth is used by a multiple access network. The most efficient use of the multiple access channel is not achieved by creating many orthogonal or nearly orthogonal subchannels; rather, it is better for every user to superimpose its signal onto the shared channel, and for a decorrelator to determine the unique hypothesis that would give rise to the measured superposition of signals. The controlled redundancy of forward error correcting codes facilitates this operation, achieving a desired error rate while drastically improving spectral efficiency. Conceptually, it may be easiest to think of error control coding and multiple access coding as distinct functions to be performed sequentially; in fact, fusing these two related and overlapping functions would reduce processing requirements and lower the total coding overhead.

The interplay between physical channel effects, such as dispersion, and system-level considerations also merits renewed attention in the context of optical CDMA. Dispersion has a single physical origin, the dependence of group velocity on the wavelength of light, but leads to two different impairments in optical CDMA systems. The first is due to the spectral spread of light contained inside a single wavelength channel. The individual laser pulses broaden as they propagate, which can lead to intersymbol interference as adjacent pulses begin to overlap at the receiver. Today, this effect is important mainly over long distances (e.g., hundreds of kilometers).

Over the comparatively short distances of interest in optical CDMA LANs, the second effect of dispersion, temporal skewing of components of codes contained in different wavelength channels, is more problematic. At the transmitter, chips in the same time slot but at different wavelengths will be synchronized; however, after traveling through a dispersive fiber, the chips with longer wavelengths will arrive at different times than chips of shorter wavelengths. If a chip is delayed enough to "spill over" into an adjacent time slot at the receiver, the temporal alignment of the code will be compromised, and the bit will be incorrectly decoded.

One solution to the dispersion problem may be controllable compensatory preskewing. Preskewing would ensure that the code arrived with the correct temporal alignment only at its intended destination; in fact, if such preskewing were integrated with the code design process, it might well be possible to exploit the effects of dispersion in effectively randomizing the received signal at all but the intended destination node. In very-high-speed optical CDMA networks, controlled preskewing could introduce additional latency since some deliberate delay would have to be added to specific wavelength channels. In such a privately owned LAN, the easiest solution may be to deploy dispersion-shifted, dispersion-flattened fiber, or perform dispersion compensation in the passive star coupler, to combat the temporal alignment problem altogether.

As multiple-wavelength optical CDMA further matures, tailored implementations, protocols, and standards will require careful attention. At the outset of this work we have claimed that the ability to implement network functionality (recipient specification, latency reduction, service differentiation) in the optical layer is not only desirable, but may be necessary as communication rates continue to grow at an explosive pace. The fundamentally noninteracting nature of the photon — the phenomenon that accounts for the great success of fiber optics — must be borne in mind when envisioning prospects for an intelligent optical network layer. New materials, concepts, and devices for optical routing and processing of signals must be made a reality if the rich spectrum of optical functionality is to deliver on its tremendous promise.

References

[1] B. Mukherjee, *Optical Communication Networks*, New York, McGraw–Hill, 1997.
[2] P. R. Prucnal, M. A. Santoro, and T. R. Fan, "Spread Spectrum Fiber-Optic Local Area Network Using Optical Processing," *IEEE J. Lightwave Tech.*, vol. LT-4, no. 5, May 1986, pp. 547–54.
[3] J. A. Salehi, "Code Division Multiple-Access Techniques in Optical Fiber Networks — Part I: Fundamental Principles," *IEEE Trans. Commun.*, vol. 37, no. 8, Aug. 1989, pp. 824–33.
[4] J. A. Salehi and C. A. Brackett, "Code Division Multiple-Access Techniques in Optical Fiber Networks — Part II: Systems Performance Analysis," *IEEE Trans. Commun.*, vol. 37, no. 8, Aug. 1989, pp. 834–42.
[5] P. A. Perrier and P. R. Prucnal, "Wavelength-Division Integration of Services in Fiber-Optic Networks," *Int. J. Digital Analog Cabled Sys.*, vol. 1, Jan. 1988, pp. 149–57.
[6] G.-C. Yang and W. C. Kwong, "Performance Comparison of Multiwavelength CDMA and WDMA + CDMA for Fiber-Optic Networks," *IEEE Trans. Commun.*, vol. 45, no. 11, Nov. 1997, pp. 1426–34.
[7] L. Tancevski and I. Andonovic, "Hybrid Wavelength Hopping/Time Spreading Schemes for Use in Massive Optical Networks with Increased Security," *J. Lightwave Tech.*, vol. 14, no. 12, Dec. 1996, pp. 2636–47.
[8] W. C. Kwong and G.-C. Yang, "Image Transmission in Multicore-Fiber Code-Division Multiple-Access Networks," *IEEE Commun. Lett.*, vol. 2, no. 10, Oct. 1998, pp. 285–87.
[9] G.-C. Yang, "Variable-Weight Optical Orthogonal Codes for CDMA Networks with Multiple Performance Requirements," *IEEE Trans. Commun.*, vol. 44, no. 1, Jan. 1996, pp. 47–55.

Biographies

ANDREW STOK (andrew.stok@utoronto.ca) received a B.Sc. degree in engineering physics from Queen's University, Kingston, Canada in 1998. In March 2000 he received an M.A.Sc. degree in photonics from the Department of Electrical and Computer Engineering at the University of Toronto. He is currently pursuing a Ph. D. degree in photonics from the same department. His research interests include the theory, design, and development of future optical networks, with a focus on optical CDMA.

EDWARD H. SARGENT holds the Nortel Junior Chair in Emerging Technologies in the Department of Electrical and Computer Engineering at the University of Toronto. He leads a team of 16 graduate and post-doctoral researchers in multiple-access optical networks, semiconductor quantum electronic devices, photonic crystal applications, and novel photonic materials. In 1999 he was awarded the Silver Medal of the Natural Sciences and Engineering Research Council of Canada and the Premier's Research Excellence Award of the Province of Ontario.