Optically Assisted Internet Routing Using Arrays of Novel Dynamically Reconfigurable FBG-Based Correlators

Michelle C. Hauer, *Student Member, IEEE, Student Member, OSA*, John E. McGeehan, *Student Member, IEEE*, Saurabh Kumar, *Student Member, IEEE*, Joseph D. Touch, *Senior Member, IEEE*, Joseph Bannister, *Senior Member, IEEE*, Edward R. Lyons, *Member, IEEE*, C. H. Lin, A. A. Au, H. P. Lee, *Member, IEEE*, Dmitry S. Starodubov, *Member, IEEE*, and Alan Eli Willner, *Fellow, IEEE, Fellow, OSA*

***Abstract*—As routing tables in core Internet routers grow to exceed 100 000 entries, it is becoming essential to develop methods to reduce the lookup time required to forward packets toward their destinations. In this paper, we employ a bank of novel thermally tuned fiber-Bragg-grating-based optical correlators to construct an "optical bypass" to accelerate conventional electronic Internet routers. The correlators are configured as a routing table cache that can quickly determine the destination port for a fraction of the incoming traffic by examining only a subset of the bits in an IP packet's 32-bit destination address.We also demonstrate a novel multiwavelength correlator based on fiber Bragg grating that can simultaneously recognize the header bits on multiple wavelengths for use in wavelength-division-multiplexed (WDM) systems. Using the optical bypass, routing table lookup times are reduced by an order of magnitude from microseconds to nanoseconds and are limited only by the speed of the optical switch.**

***Index Terms*—Optical communications, optical correlators, optical signal processing, wavelength-division-multiplexed (WDM) networks.**

I. Introduction

In present-day fiber-optic networks, data packets are converted to electrical form at each node to process their headers and make routing decisions, as shown in Fig. 1(a). As routing tables grow in size, more memory accesses are required to determine the next-hop address and appropriate output port to which to forward each packet. The associated increase in routing-table lookup times is becoming a significant source of latency in the network core. To make matters worse, the transmission capacity of optical fibers is rapidly increasing, forcing the routers to accommodate more packets, more often. Since routing tables will continue to grow and transmission rates are always on the rise, it is essential to develop methods to reduce the lookup time required to forward incoming IP packets. This motivates the need for optical designers to examine the operations of current electronic routers to determine where it may be feasible to employ optical techniques to assist the electronics in making ultrafast routing decisions.

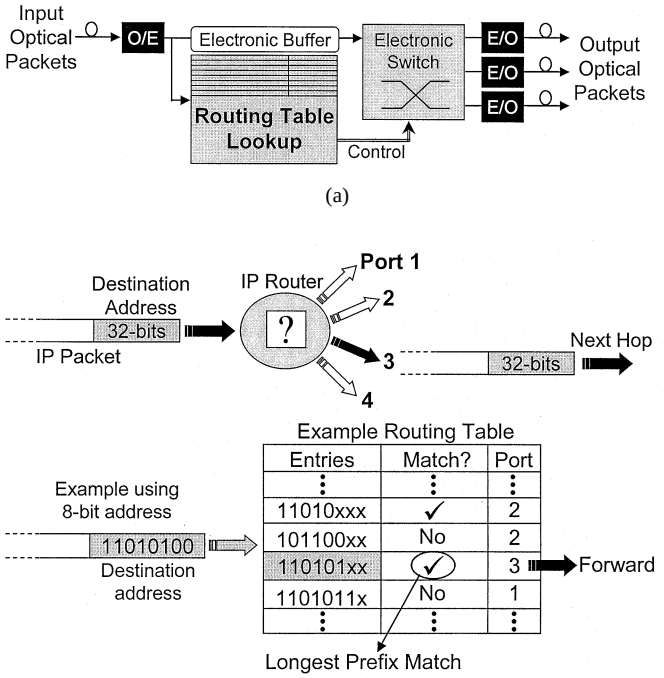

Fig. 1. (a) Electronic Internet router architecture. Incoming optical packets are converted to electrical form and stored in a buffer while their 32-bit addresses are compared to a large routing table to determine to which output ports to forward them. The packets are electronically switched and converted back to optical form for transmission. (b) Longest prefix match example using an 8-bit address.

IP routers perform two primary functions: routing and forwarding. *Routing* is the process of generating the lookup table of destination addresses and corresponding next-hop addresses and output ports. Routing tables in the network core are fairly stable and are typically recomputed on a timescale of tens of minutes to reflect the continually evolving layout of the network. *Forwarding* is the process of steering packets toward their

Manuscript received December 20, 2002; revised May 22, 2003. This work was supported in part by Intel and in part by the NSF Optical Networks initiative.

M. C. Hauer, J. E. McGeehan, S. Kumar, and A. E. Willner are with the Department of Electrical Engineering—Systems, University of Southern California, Los Angeles, CA 90089-2565 USA (e-mail: hauer@usc.edu).

J. D. Touch and J. Bannister are with the Information Sciences Institute, University of Southern California, Marina del Rey, CA 90292-6695 USA (e-mail: touch@isi.edu).

E. R. Lyons is with Raytheon, El Segundo, CA 90245 USA (e-mail: erlyons@raytheon.com).

C. H. Lin, A. A. Au, and H. P. Lee are with the Department of Electrical and Computer Engineering, University of California Irvine, Irvine, CA 92697 USA (e-mail: hplee@uci.edu).

D. S. Starodubov is with Sabeus Photonics, Chatsworth, CA 91311 USA (e-mail: starodubov@sabeus.com).

Digital Object Identifier 10.1109/JLT.2003.819144

destinations by comparing their 32-bit destination addresses to entries in a routing table using a longest prefix matching algorithm. Thus, routing is analogous to drawing a map and forwarding is the act of following its directions. The basic concept of forwarding is illustrated in Fig. 1(b) for a packet with a simple 8-bit destination address of "11 010 100." This address is compared to entries in the routing table, where two matches are found. The matching entry with the longest prefix is chosen and directs the packet to output port 3.

The forwarding process can be time consuming given that core routing tables have grown to contain more than 100 000 entries. Lookup times are presently on the order of microseconds. Given that a significant portion of Internet packets entering routers are short, 40-byte TCP/IP acknowledgement packets, nanosecond lookup times are needed to achieve the desired terabit/second throughput. Some efforts are being made to improve electronic hardware architectures and search algorithms to reduce lookup times [1], but the ideal case would be for the packet headers to be processed on-the-fly using optical signal-processing techniques, so that the only limitation to throughput is the speed of the optical switch (currently a few nanoseconds). A true all-optical router would therefore need to be capable of 24-bit lookups into 100 000-entry tables at $\geq$ 10 Gb/s (since only 24 of the 32 address bits are generally significant for packet-forwarding in the network core). Unfortunately, such capabilities are beyond current optical technologies, which are limited to optical correlation techniques. These techniques inherently suffer from high optical splitting losses that limit them to matching incoming packets against a few patterns that are a few bits long. However, some recent developments hint at the feasibility of a partial solution in which presently achievable optical correlators may be combined with a novel routing-table optimization algorithm to reduce lookup times by at least an order of magnitude from microseconds to nanoseconds.

Given that most core routers have only four to eight outgoing ports, it may be possible to determine a packet's output port by looking at only a small subset of the 24 bits in the destination address. So although building a true all-optical router is beyond current optical technologies, it is feasible to build an "optical bypass" to accelerate a conventional router. A subset of the traffic would be routed by the optical bypass without any O/E conversion, at increased throughput and decreased latency. The remainder of the traffic, which requires more complicated processing, is handled by a conventional electronic router. Previous research has shown that by utilizing a subset of the routing table with as few as 100 of the most popular entries, the router can still successfully forward as much as 90% of the incoming traffic [2]. The remaining challenge is to determine how to design a 24-bit input, 100-entry optical index using a manageable number of optical correlators that interrogate only a small subset of the bits in the destination address.

The optical correlators required for this application must be tunable and designed to easily scale to 40 Gb/s and beyond. Optical correlators are typically implemented with tapped-delay-line structures that split the optical signal into several branches, where each successive branch delays the signal 1 bit-time longer than the previous branch. The tiny distances required to achieve these differential 1-bit delays in fiber at bit rates $>$ 10 Gb/s present a serious challenge for previously reported correlator designs. These designs include optical splitters followed by fiber delay lines terminated with fiber mirrors [3] and arrays of discrete fiber Bragg gratings (FBGs) tuned with separate piezoelectric stretchers [4]. For these cases, the intermirror or intergrating spacings would have to be 2.5 mm for a 1-bit round-trip time delay at 40 Gb/s—an impractical length for devices using discrete fiber components.

In this paper, we propose and experimentally demonstrate a correlator design in which an FBG array is constructed from a single uniform fiber grating that is divided into separate, electrically tunable sections using thin-film microheaters. The precision of the heater deposition defines the spacing between gratings. Since thin-film heaters can be fabricated with lithographic precision, spacings down to hundreds of micrometers are achievable with this technology, enabling the correlator to readily scale to higher bit rates.

An additional advance over previously reported correlator architectures is that our approach uses two grating arrays per correlator: one to correlate with the desired "1" bits in the input signal and another to match with the desired "0" bits. We show that the combined outputs from both of these correlators are necessary to produce unique correlation outputs for all 2^n possible n-bit sequences.

In wavelength-division-multiplexed (WDM) systems, packets arrive at routing nodes on multiple wavelengths. For this case, a separate bank of correlators is required for each incoming wavelength channel. To reduce the number of components required in a WDM routing node, we demonstrate a correlator design that is constructed with sampled FBGs that enables the headers on multiple wavelengths to be simultaneously tested against a stored bit pattern. The correlation operation can therefore be performed prior to demultiplexing and switching. This eliminates the need for multiple sets of correlators and significantly reduces the required number of components in a WDM routing node.

This paper is structured as follows. In Section II, the concept of an optically assisted Internet router is described along with the constraints for the algorithm needed to define and configure the optical routing table. Section III provides an overview of digital optical correlation. Section IV presents a novel FBG-based correlator design and includes an experimental demonstration in which the correlators are used to process and switch 10-Gb/s optical packets. In Section V, a multiwavelength FBG-based optical correlator is described and demonstrated in a WDM system. Section VI includes a discussion of important fabrication and performance considerations associated with using FBG arrays for optical correlation. This paper is summarized in Section VII.

II. Optically Assisted Internet Router

To implement an effective optical bypass for an electronic router, the key design decision is to combine a software algorithm with a small set of dynamically configurable FBG-based optical correlators. A conceptual diagram showing how the optical bypass is implemented in an IP router is shown in Fig. 2(a).

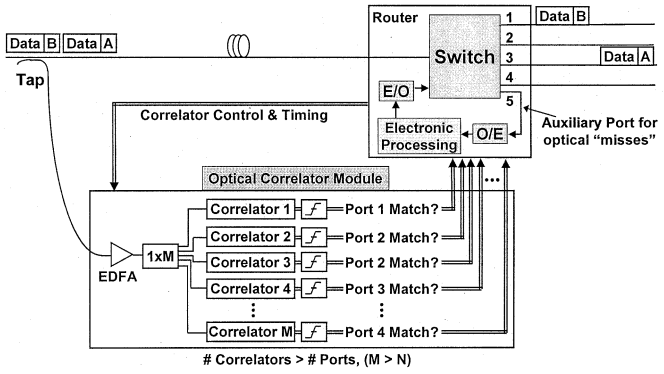

(a)

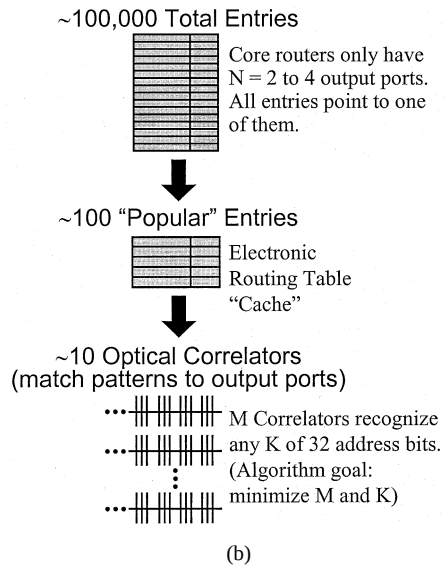

(b)

Fig. 2. (a) Conceptual diagram of an Internet core router with N output ports (typically two to four) assisted by a bank of M optical FBG-based correlators that are dynamically configured by a software algorithm. (b) The algorithm searches the port can be determined by examining only a subset K, of the n bits in the destination address. The correlators are then configured corresponding to these groups of entries. If all the correlators fail to match the incoming packet address, then the packet is switched to an auxiliary port and processed by conventional electronics.

A small portion of the incoming optical packet stream is tapped off and sent to the correlator module. The optical signal is amplified, split, and sent to M correlators, each of which can be configured to produce a "match" signal for any number K of the 24 significant bits in the destination address.

Previously, we reported on an algorithm for generating a routing table cache containing as few as 100 entries that could successfully route as much as 90% of the traffic entering an IP router [2]. Since a 100-entry lookup table is still too large implement optically, the goal here is to define an algorithm to represent this 100-entry table with as few as eight or ten optical correlators. The fact that Internet routers typically have only two to four outgoing ports makes this feasible. Even if there are 100 000 entries, each of them only points to one of two or four possible destinations. The algorithm searches for patterns among these entries for which the output port can be determined by examining only a subset K of the n bits in the destination address, with K ideally less than five. An additional goal of the software is to determine the optimum groupings of these entries so as to minimize the number M of optical correlators required. The software then configures each optical correlator to represent one of the resulting groups.

This process is illustrated in Fig. 2(b). Note that there will likely be multiple correlators corresponding to each output port since there will be groups of bit patterns that the algorithm determines should all be routed to a particular port. For example, it is feasible that two correlators, each of which are configured to match a different subset of address bits, both route packets to the same port when they get a match. The algorithm must recompute the optical lookup table and reconfigure the correlators each time the routing table is updated. Threshold detectors are used at the outputs of the optical correlators to provide an electrical match/no-match signal to the optical switch. The switch uses these signals to determine to which output port each packet should be forwarded. If the correlators fail to find a match for an incoming packet, the switch routes the packet to an auxiliary port for electrical processing by a conventional electronic router [see port 5 in Fig. 2(a)]. For example, if the algorithm predetermines that any incoming packets with bit positions 1, 4, and 5 equal to 1, 0, and 1 should go to port 1, then a correlator is configured to provide a "match" signal for any input pattern with "1xx01" for its first five bits and anything else for the rest of the address (where the "x"s indicate "don't care" bits that can be either a "1" or a "0").

While the optical bypass enables on-the-fly forwarding of incoming packets, there are some issues associated with forwarding packets without converting them to electronics to process and update their headers. For example, the IP header's time-to-live (TTL) field is not decremented and the header checksum is not recomputed, whereas protocol requires that both of these operations occur at each network hop. One potential solution to this problem is to revise the protocol to allow for packets to traverse a small number of core network hops without O/E conversion and then update these fields once they reach a fully electronic router at the core edge [1616]. Alternatively, some advanced optical signal-processing techniques are being developed to directly operate on these fields in the optical domain, though this research is still in the early development stages [5].

The constraints for the algorithm can be formalized as follows. If we have a router with N ports and a bank of M correlators, we can use the correlator bank as a routing cache, which stores a fraction of the routing table's commonly accessed n-bit addresses. A formulation of this problem is as follows. Let $F : \{0,1\}^n \rightarrow \{1,2,\ldots,N\}$ denote the routing table, which is a function that maps any destination address to its output port. The ith correlator is configured to recognize a given bit pattern, where the pattern consists of an n-bit string of 0s, 1s, and "don't care" bits, i.e., its pattern specification is the set of vectors $B_i = B_{i1} \times B_{i2} \times \cdots \times B_{in}$, where $B_{ij} \in \{\{0\},\{1\},\{0,1\}\}$ for $1 \leq i \leq M$ and $1 \leq j \leq n$. The correlator is inherently limited and can recognize at most K bits of an address, so that at least $n - K$ bits of the address are "don't care" bits. There are three formal requirements to satisfy when implementing a routing cache from the correlator bank.

1) No false positives: if $d_1 \in B_i$ and $d_2 \in B_i$, then $F(d_1) = F(d_2)$.
2) Correlator limitations: minimize M and K.
3) Nonoverlapping correlators: if $i \neq j$, then $B_i \cap B_j = \emptyset$.

The set of pattern specifications should be efficiently computable from the routing table F, and the correlators should provide a sufficiently high hit rate for traffic entering the router. In today's networks, typical values for the parameters of the problem are $N = 4$ (the number of output ports in a typical Internet router), $M = 8$, $n = 24$ [the most significant bits of an Internet Protocol Version 4 32-bit address], and $K = 5$ (a constraint of the optical technology). The primary focus of this paper is to demonstrate the feasibility of implementing a bank of optical correlators that will meet the requirements above.

III. Digital Optical Correlation

A. Optical Tapped-Delay-Line Correlators

Correlation, or matched filtering, is an important signal-processing function. The purpose of a correlator is to compare an incoming signal with one that is "stored" in the correlator. At the appropriate sample time, a maximum autocorrelation peak will be produced if the input signal is an exact match to the stored one. This function has historically been used in conjunction with special coding techniques to pick a desired signal out of noise, an essential requirement for RADAR and code-division multiple-access (CDMA) systems. In optical communication systems, correlators can likewise be used to recognize particular bit patterns, enabling applications such as optical CDMA networks and header recognition (the focus of this paper). The aim of this section is to give a brief overview of the concepts that apply to the correlation of digital binary waveforms that are intensity-modulated onto optical carriers.

A common implementation of an optical correlator is the tapped delay line. A basic optical tapped delay line correlator is shown in Fig. 3(a). The delay line is configured to match the correlation sequence "1101." Thus, the delay line requires four taps (one for each bit in the desired sequence), weighted by the factors 1, 1, 0, and 1, respectively. The weights are implemented by placing a switch in each path that is closed for weight $= 1$ and opened for weight $= 0$. The incoming optical bitstream is equally split among the four fiber-optic delay lines. Each successive delay line adds one additional bit of delay to the incoming signal before the combiner, where the powers of the four signals are added to yield the correlation output function. This function is sampled at the optimum time T_s and passed through a threshold detector that is set to detect a power level above two, since the autocorrelation peak of "1101" with itself equals three (or, more specifically, three times the power in each "1" bit). The output is detected using a photoreceiver, and a simple electronic decision circuit is used to compare the correlation output to the threshold value. The high-speed advantage of optics still prevails in this case since the correlation function is produced in the time it takes the signal to traverse the optical correlator. The decision circuit only needs to be triggered at the packet rate, which is often in the range of kHz to MHz, depending on the number of bits in each data packet. For example, for a stream of short, 40-byte packets (320 bits/packet) at 40 Gb/s, the packet rate is only 125 MHz. The mathematical function describing the tapped-delay-line correlator is

$$y(t) = \sum_{j=0}^{n-1} x(t - jT_{\text{bit}})h(jT_{\text{bit}}) \tag{1}$$

where n is the number of bits in the correlation sequence, T_{bit} is one bit period, $x(t - jT_{\text{bit}})$ is the input signal delayed by j bit times, and $h(jT_{\text{bit}})$ represents the j weights that multiply each of the j-bit delayed input signals. For a phase-modulated system, the same operation is performed by replacing the switches with the appropriate optical phase-shifters to match the desired bit pattern (e.g., "$\pi\pi 0\pi$" instead of "1101"). Fig. 3(b) illustrates the delay-and-add operation of the correlator for the case when the three 4-bit words "1011," "1101," and "0101" are input to the correlator, where the second word is an exact match to the desired sequence. Since the correlation function for two 4-bit words is 7 bits long and the peak occurs during the fourth time slot, the correlation output is sampled every four bits and compared to a threshold as shown in Fig. 3(c). As expected, the correlation output for the second word exceeds the threshold, while the first and third samples produce no matches. Note that for an input signal L bits long, the length of the correlation output will be $L + n-1$ bits long.

B. The Need for "Ones" and "Zeros" Correlators

Note that the correlator as shown in Fig. 3(a) will also produce a level "3" peak that is above the threshold at time T_s for a "1111" input, which is not the desired bit pattern. This is because the open switch in the third delay line, corresponding to the third correlation bit, does not "care" if the third bit is a "1" or a "0" since it does not pass any light. Thus, the correlator as shown is really configured to produce a match for the sequence "11x1" where the "x" indicates a "don't care" bit that can either be "0" or "1."

This is not an issue in optical CDMA systems, where the set of codewords can be specifically designed to maintain a constant number of "1" bits in each codeword. However, for our header recognition application, which must be able to uniquely recognize any of the 2^n possible n-bit sequences, this situation will result in "false-positive" matches whenever a "1" is present where a "0" bit is desired. To overcome this problem, we add a second correlator that is configured in complement to the first one and produces a "match" signal when *zero* power is present at the sample time. This is accomplished by placing a NOT gate at the output of the threshold detector that is set just above level zero. If the power goes above the threshold, this indicates that at least one "1" bit is present where a "0" is desired, and the NOT gate will convert the high output to a low one to indicate "no match" for this correlator. This correlator therefore correlates with the desired "0" bits in the sequence and is thus called a "zeros" (or "0s") correlator. Likewise, the originally described correlator is called a "ones" (or "1s") correlator. In the zeros correlator, the switches are closed for desired "0" bits and open otherwise (or, as explained in the next section, the FBG mirrors reflect for desired "0" bits and are tuned away otherwise). Thus, the "1" bits are "don't care" bits in a zeros correlator. By

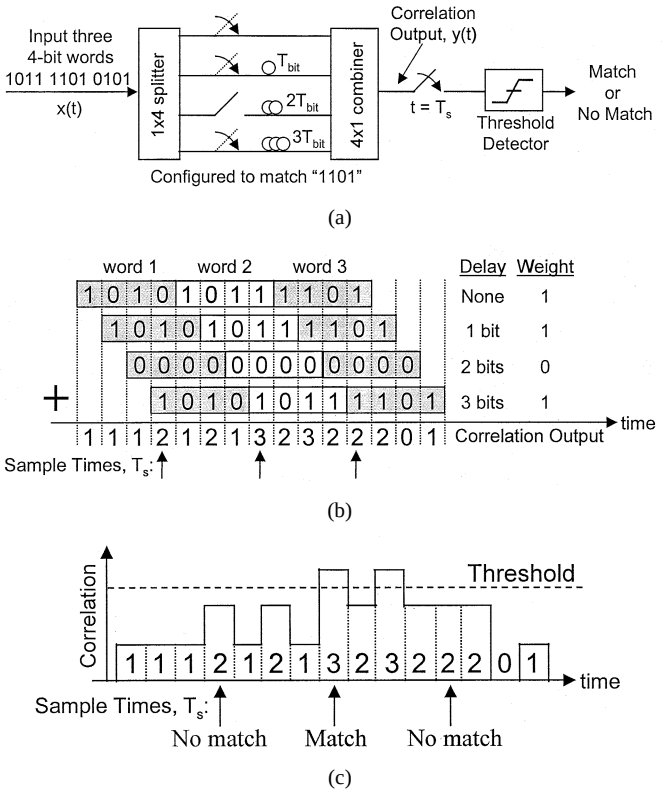

Fig. 3. (a) A basic implementation of a fiber-optic tapped-delay-line correlator configured to produce an autocorrelation peak at the sample time, T_s, for the sequence "1101." (b) The weighted delay-and-add computation of the digital correlation output for three input words when correlated with the bit pattern "1101." The three optimum sample times are labeled. (c) An intensity profile of the correlation output from this tapped-delay-line correlator. Only when the intensity is above the threshold at the sample time is a match signal produced.

combining the photodetected outputs of the ones and zeros correlators with an electronic AND gate, a final "match" signal will only be produced when the input pattern uniquely matches the desired correlation sequence.

An illustration of how the combination of ones and zeros correlators can avoid false positive matches is depicted in Fig. 4. The desired correlation sequence is "1001," meaning the ones correlator is configured to match a "1xx1" pattern and the zeros correlator will produce a match for an "x00x" pattern. In Fig. 4(a), the incoming sequence is "1001," and so the ones and zeros correlators both produce "match" signals, resulting in a "match" signal at the output of the AND gate. In Fig. 4(b), the input sequence is a "1101," causing the ones correlator to still produce a "match" signal (this would be a false positive match if only this correlator were present), while in the zeros correlator, the undesired "1" bit in the second time-slot of the input causes the power at the sample time to exceed the threshold, resulting in a "no match." The combination of the "match" and "no match" signals in the AND gate produces the correct "no match" result.

IV. Reconfigurable FBG-Based Optical Correlator

The optical tapped-delay-line structure shown in Fig. 3(a) requires a separate fiber branch and an optical switch for each bit in the desired bit pattern, making it impractical to construct a bank of 24-bit correlators. Moreover, the length of each fiber

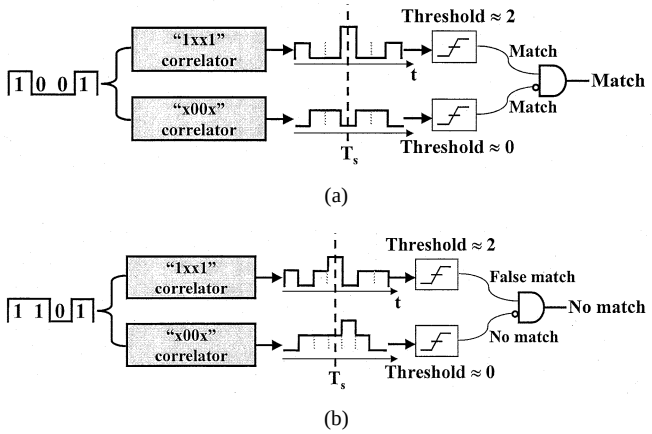

Fig. 4. The concept of combining "ones" and "zeros" correlators to uniquely recognize a bit sequence and avoid false-positive matches. The "ones" correlator tests for "1" bits in the correlation sequence and the "zeros" correlator tests for "0" bits. The correlators shown are configured to recognize a "1001" pattern when their outputs are combined in an AND gate. (a) The input pattern "1001" results in a match for both correlators, producing a final "match" decision at the output. (b) The input pattern "1101" results in a match for the "ones" correlators but a "no match" for the "zeros" correlator. The combination of these two outputs produces the correct "no match" decision at the output of the AND gate.

branch must be cut to precisely provide the requisite 1-bit delays between successive branches. This corresponds to a differential length of 2 cm of fiber at 10 Gb/s and 5 mm at 40 Gb/s. A simpler, more producible, and manageable correlator may be constructed by writing a series of fiber Bragg grating mirrors into a single length of fiber, as shown in Fig. 5. In this case, the reflectivities of the FBG mirrors provide the same weighting function as the optical switches in Fig. 3. The gratings representing desired "1" bits are tuned to reflect (closed switch case) while those representing "0" bits are tuned to be transparent (open switch case). Since the light makes a double pass through the array, the spacing between FBGs must correspond to 1/2 of a bit time to produce a round-trip delay of 1-bit-time. An optical circulator is placed at the input to route the counterpropagating correlation output to the threshold detector. Aside from these differences, the operation of the correlator is identical to that described for Fig. 3.

A nice feature of FBG filters is that their reflection spectra can be tuned by either heating or stretching the fiber. For FBGs fabricated to act as high-reflectivity mirrors, the reflectivity is nearly 100% at the center of the reflection spectrum and falls off toward zero outside the grating bandwidth. By tuning the FBG so that the signal wavelength intersects with the rising or falling edge of the filter's passband, the reflected energy at that wavelength will be reduced from 100% toward 0% as the grating is tuned. The array in Fig. 5 is a 4-bit ones correlator, configured to match the pattern "11x1." As shown in the figure, this is accomplished by tuning the reflectivities of the three FBGs representing "1" bits to 23%, 38%, and 100%. These reflectivities may be computed using the following recursive equation:

$$\frac{R_n}{(1-R_n)^2} = R_{n+1}. \tag{2}$$

This equation is derived by requiring that the light reflecting off each FBG have equal power when it exits the correlator.

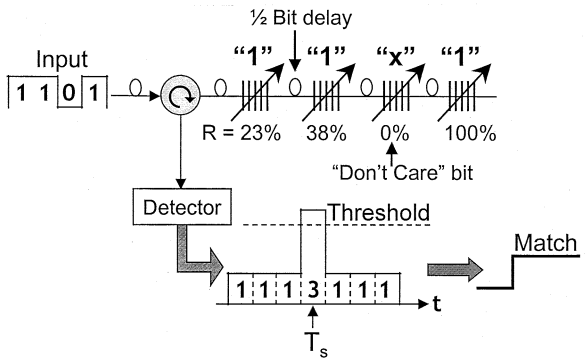

Fig. 5. Illustration of a 4-bit optical correlator implemented with an array of FBGs. The FBGs representing "1" bits are tuned to partially reflect, while the "0" or "don't-care" FBGs are tuned to fully transmit. The delay between gratings is 1/2 of a bit-time to provide a round-trip delay of 1 bit-time. The correlator shown is configured to recognize a "11x1" pattern, where the "x" can either be a "1" or "0" bit.

Thus, the reflectivity of the last grating in the array should be set equal to one, and then the reflectivities of all the preceding gratings can be calculated using this recursive equation. For the case shown in Fig. 5, with $R_4 = 1$ and $R_3 = 0$ ("don't care" bit), we get $R_2 = (1 - R_2)^2$, resulting in $R_2 = 0.38$. Then, $R_1 = 0.38(1-R_1)^2$, yielding $R_1 = 0.23$. The reflectivity of the first FBG is fairly low since the portion of light reflected from it must equal the optical power that makes a double pass through all the other gratings in the series. This lossiness places a limit on how many bits can realistically be in the correlation sequence and is the reason that our software algorithm aims to locate addresses in the routing table that can be forwarded by examining only $K = 5$ or fewer of the 24 significant address bits. With five desired "1" bits in a correlation sequence, the reflectivity of the first reflective FBG is 12.4%. Since the autocorrelation peak will contain five stacked optical bits, the peak correlator output power will be $(5)(0.124)P_{\text{input}}$, where P_{input} is the optical power in an input "1" bit. This power must be greater than the noise floor at the correlator output in order for the threshold detection to work properly and places a constraint on P_{input}.

In practice, the necessary reflectivities of the gratings are determined by repeatedly sending a single pulse into the FBG array and observing the powers of the multiple, time-delayed output pulses on an oscilloscope. The gratings are tuned until the output pulses have equal power. When the reflectivities are properly tuned, the response of the array to a single input pulse is the bit pattern that the correlator is configured to recognize.

The FBG correlator shown in Fig. 5 uses a discrete array of FBGs, in which each grating is written separately and great care must be taken to ensure that the spacing between gratings precisely equals one-half of a bit time. This becomes increasingly difficult at higher bit rates, where the center-to-center spacings range from 1 cm at 10 Gb/s down to 1 mm at 100 Gb/s. To construct an electrically tunable correlator that more readily scales to higher bit rates, we chose to deposit a series of thin-film microheaters onto the surface of a single, long uniform grating as shown in Fig. 6(a). This way, the placement of the heaters defines the grating spacings. When a voltage is applied to one of the heaters, the effective index of the portion of the grating directly beneath the heater varies in response to the rise in temperature. This causes the filter spectrum of that small section to shift toward longer wavelengths as shown in Fig. 6(b). This design simplifies the correlator construction since uniform FBGs that are several centimeters in length are simple and inexpensive to fabricate. The metal heaters can be deposited with simple e-beam evaporation, or, to achieve very close spacings, by integrating the heater array onto a silicon substrate using conventional lithographic techniques and then affixing the fiber to the heater array [14]. The modulation depth of the grating's effective index δn_{eff} should be large ($\sim 10^{-3}$) to achieve a strong grating with high reflectivity to ensure that the shorter subgratings will still have high peak reflectivities.

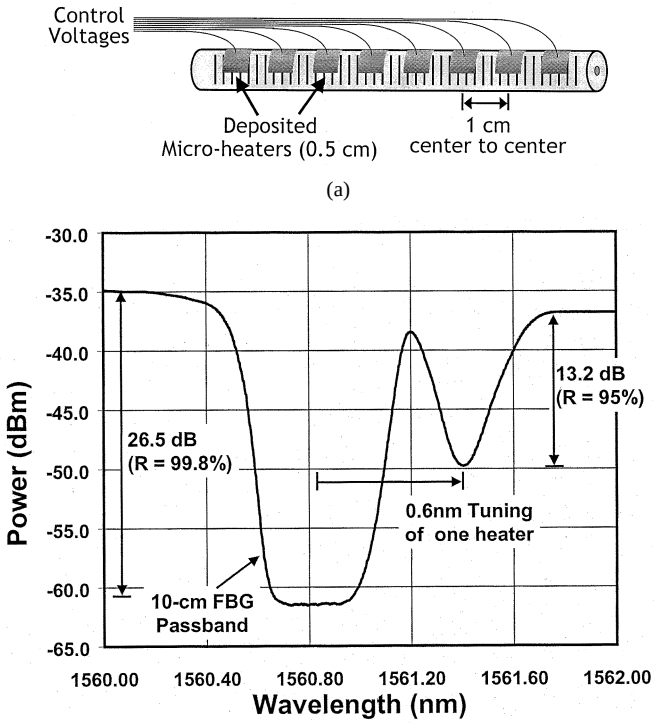

Fig. 6. (a) A 10-cm-long FBG with eight deposited metallic thin-film heaters to create a tunable array of subgratings. (b) Transmission spectrum of the FBG after a voltage has been applied to one of the heaters to tune its corresponding subgrating by 0.6 nm from the central passband of the 10-cm-long FBG. The peak reflectivity of the subgrating is ~95%.

This design can accommodate the reconfigurable 24-bit correlator needed for our optical bypass. The long FBG is fabricated such that its reflection peak is out-of-band of the incoming wavelength, i.e., the FBG is transparent at the incoming signal wavelength unless one of the heaters is tuned. This is an efficient design because at any given time, the correlator will be configured to recognize only a small subset of the 24 address bits, meaning most of the bits will be "don't care" bits, and for these bits the heaters will simply remain "off." Only the few heaters representing "1" bits in the ones correlator and "0" bits in the zeros correlator need to be tuned.

To demonstrate the feasibility of this design, we constructed two 8-bit correlators: one to operate as the ones correlator and the other for the zeros. The correlators were fabricated using e-beam evaporation to deposit an array of eight metallic heaters onto two 10-cm-long uniform fiber Bragg gratings at a center wavelength of ~1560.8 nm. The heaters are 0.5 cm long with

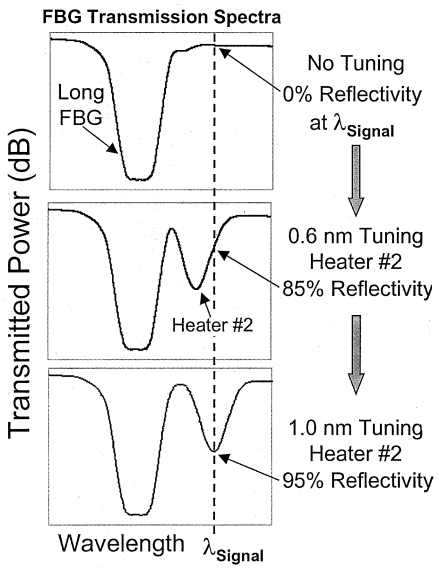

Fig. 7. Tuning the reflectivity at the signal wavelength from 0% to 95% by tuning the second heater so that the signal wavelength intersects with different points along the edge of the subgrating passband.

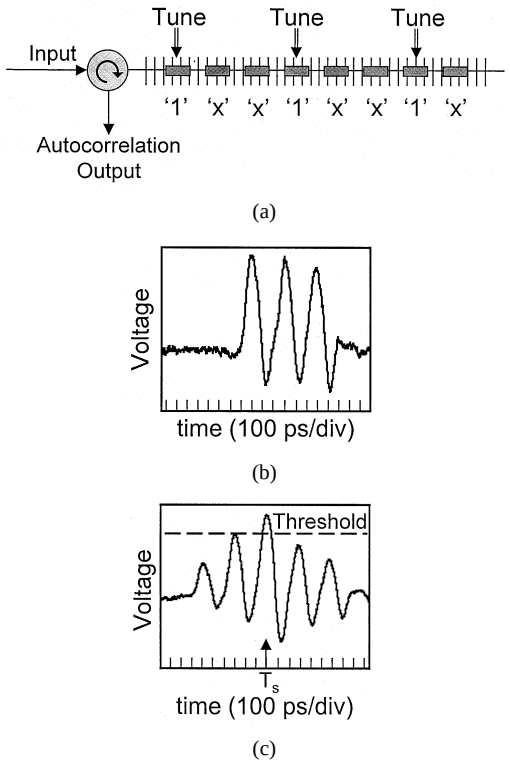

Fig. 8. (a) A 10-cm-long FBG with eight microheaters configured as a ones correlator to recognize the 10-Gb/s pattern "1xx1xx1x." (b) Oscilloscope trace of the correlator output for a single input pulse showing that the reflectivities are properly tuned to provide equal-power output pulses at the first, fourth, and seventh bit positions. (c) Autocorrelation output "010 020 030 020 010" when the input is "10 010 010." Since the input matches the correlator configuration, the central peak exceeds the threshold at the sample time.

center-to-center spacings of 1 cm, corresponding to 1/2 of a bit time in fiber at 10 Gb/s. The heaters consist of a 15-nm titanium layer for good adhesion to the glass and a 120-nm layer of gold for good electrical conductivity.

The transmission spectrum of the grating is shown in Fig. 6(b) with one of the heater subgratings tuned by 0.6 nm from the center of the main FBG's passband. The reflectivity of the main passband is nearly 100% since it is 10 cm long, and the peak reflectivity of the shorter subgrating is approximately 95%. Fig. 7 shows how the reflection of the signal light varies from 0% to 95% as the second heater is tuned. The input signal wavelength was set to 1561.8 nm, requiring 1 nm of tuning to obtain maximum reflection. This 1-nm shift required $\sim$25 mW electrical power applied to the heater, yielding an FBG tuning efficiency of $\sim$40 nm/W. We can also estimate that a temperature rise of about 70 °C is required to produce the 1 nm of tuning using a typical FBG sensitivity of about 14 pm/°C [6]. Assuming room-temperature operation, this means that a constant local temperature of approximately 100 °C is required for each of the tuned FBG sections. An important practical consideration for the operation of the FBG-based optical correlator will be the long-term thermal stability of the FBG at its maximum operating temperature. It is well known that the reflectivity of FBGs can decrease over time, especially at elevated temperatures. To ensure the long-term stability of the gratings in a practical design, it will be important to anneal the FBGs at a temperature well beyond the expected maximum operating temperature [7]. One interesting possibility with the evaporated microheaters would be to use the heaters themselves to perform the grating annealing as part of the correlator manufacturing process.

An example correlator configuration for examining 3 of 8 bits is shown in Fig. 8(a). It is configured as a ones correlator to recognize the 10-Gb/s pattern "1xx1xx1x" by applying voltages to the first, fourth, and seventh heaters to tune these portions of the FBG to reflect at the input wavelength. Fig. 8(b) shows an oscilloscope trace of the correlator response to a single input pulse, which is simply the stored correlation sequence "10 010 010." This response is used to adjust the heater voltages until the three "1" bits are equal height, indicating that the subgrating reflectivites are properly tuned. Fig. 8(c) shows the measured trace of the autocorrelation output "010020030020010" that results for any input of "1xx1xx1x," where the "x" bits can be either "1s" or "0s." Since the input matches the correlator configuration, the central peak exceeds the threshold at the sample time.

The experimental setup used to demonstrate the optical bypass at 10 Gb/s is shown in Fig. 9. The incoming nonreturn-to-zero (NRZ) data packets are 53 bytes long with 8-bit headers and a 6.4-ns guard time between them to accommodate the switching time of our $LiNbO_3$ optical switch. The threshold of the ones correlator decision circuit is set to detect two stacked bits, whereas the zeros correlator threshold is always set just above zero. A packet-rate clock signal ($\sim$21 MHz) is used to trigger the decision circuits to sample the correlator outputs at the proper time. In a full system implementation, this timing signal would be generated by a previous module that detects the optical packet's arrival time [8]. The two correlation outputs are combined with an AND gate that produces a high signal when there is a match and a low signal otherwise (this signal remains latched until the next sampling result). The output match signal is amplified and used to drive a two-port $LiNbO_3$ optical switch.

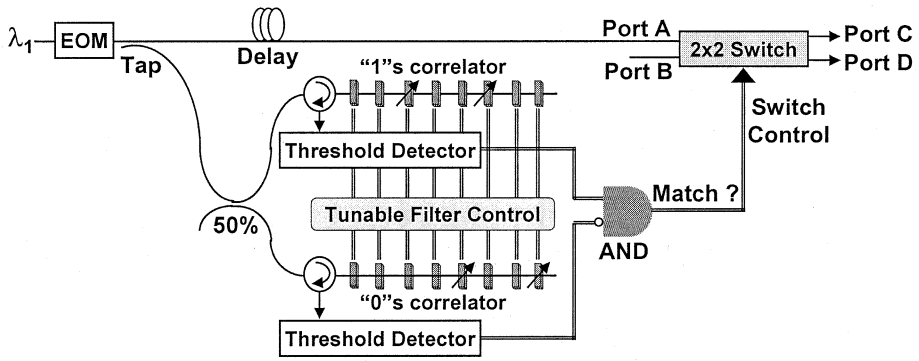

Fig. 9. Experimental setup. The FBG heaters are tuned to configure the correlator to match any "xx1x01x0" pattern. The output of the electronic AND gate drives a 2×2 $LiNbO_3$ optical switch and is set to switch packets with matching headers to port C and all others to port D.

To demonstrate the successful operation of the optical bypass function, the FBG correlators were tuned to recognize an "xx1x01x0" pattern ($K = 4$ of 8 bits), and the correlation output was used to route packets with matching headers to port C of the two-port optical switch. The experimental results showing the oscilloscope traces of the packets and switch control signals are shown in Fig. 10. The input packet stream contains four packets with different headers, the second of which matches the correlator configuration. The second row shows the matching signal at the output of the AND gate, which is amplified and used to drive the optical switch. As expected, the output goes high during the second packet. The last two lines show the successful routing of the packet with a matching header to port C and all nonmatching packets to port D.

The "lookup time" required to forward packets using these optical correlators is simply the time it takes for the light to propagate through the correlators and for the optical switch to flip, which is on the order of a few nanoseconds. Thus, the lookup time for packets that are forwarded by the optical bypass is reduced by an order of magnitude, from microseconds (for electronic lookups) to nanoseconds, and is limited only by the optical switching time. As for the reconfiguration time of the thermally tuned FBG correlators, they will have a time constant of about 1 s, which is typical of thermally tuned all-fiber devices [9]. This slow reconfiguration time should not be a problem in practice because the correlators are reconfigured only when the routing table is recomputed, which at worst occurs several times per day. However, if faster reconfiguration times are desirable, etched cladding fibers can be used to produce significantly faster tuning speeds [14].

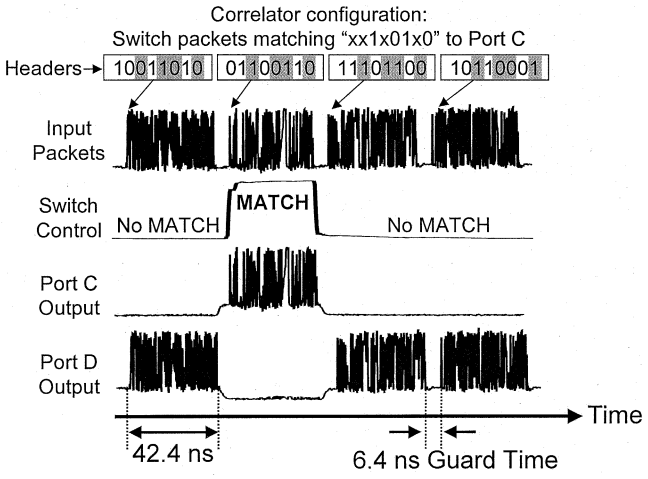

Fig. 10. Experimental results showing the successful correlation and switching of the matching packet to port C and all others to port D. The 6.4-ns guard time between packets is to allow for the switching time of the optical switch.

V. Multiwavelength FBG Correlator

The above method for optical header-subset recognition acts on a single WDM channel, thus requiring N complete modules in order to recognize the headers on N different WDM channels. A correlation module that enables reconfigurable optical correlation of multiple WDM channels simultaneously can significantly reduce the number of components required in a WDM routing node. Using a set of discrete sampled fiber Bragg gratings, a multiwavelength FBG correlator can be constructed. A sampled FBG is an FBG that has a superstructure written on top of the grating for which the Fourier transform produces a reflective time delay that is replicated at equal wavelength spacings [10]. When this type of FBG is stretched or heated, the entire reflection spectrum shifts, causing the reflectivity at each wavelength to experience the same variation. Thus, the correlation sequence can be reconfigured for all incoming channels simultaneously.

A conceptual diagram illustrating how a WDM header recognition (HR) module is employed in a routing node is shown in Fig. 11. The module may still be used as an optical bypass for an Internet router, but in this case the packets on all incoming wavelengths can simultaneously be compared to the entries in the optical lookup table. A portion of the incident light from the WDM channels is tapped off and sent into the correlation module. The module uses a bank of tunable sampled-FBG arrays for comparing the incoming packet addresses against a stored set of patterns. The correlator outputs are demultiplexed and separate decision circuits are used to sample the output for each wavelength. The resulting match/no-match signals are used to control an optical cross-connect to properly switch the packets on each wavelength.

While a sampled FBG correlator can be constructed in a manner similar to a standard FBG correlator, the spacing

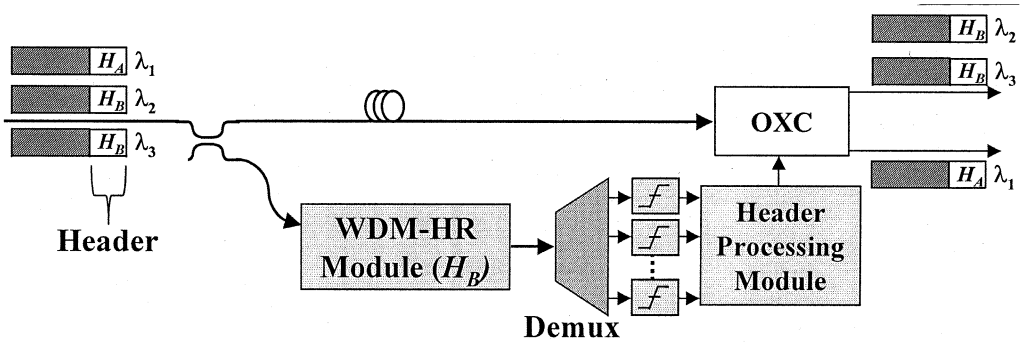

Fig. 11. Multiwavelength correlation module. HR is performed simultaneously on the WDM channels prior to demultiplexing using a single bank of WDM optical correlators. The correlation outputs are demultiplexed and threshold detected. A header processing module uses this information to control the switches in an optical cross-connect (OXC).

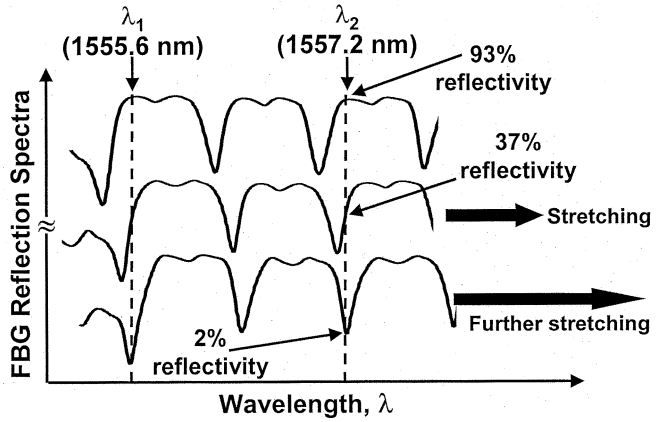

Fig. 12. Wavelength spectrum of one of the sampled FBGs used in the experiment. The reflectivity at the two input wavelengths (1555.6 and 1557.2 nm) is tuned from 93% to nearly zero by stretching the grating.

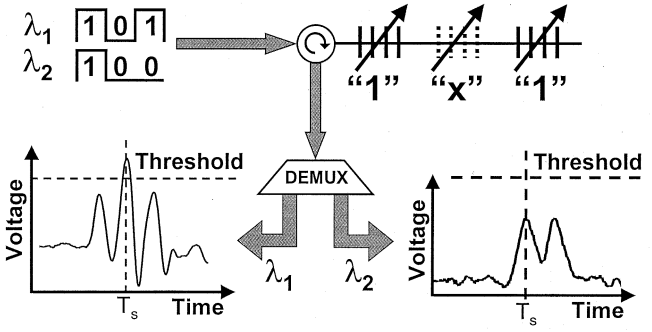

Fig. 13. Operation of 3-bit multiwavelength correlator using sampled gratings that are tuned to match a "1x1" pattern. Oscilloscope traces show the output correlation function for each wavelength after demultiplexing. The output for λ_1 is the autocorrelation response for a "101" sequence. However, the input sequence on λ_2 ("100") does not match the correlator and produces a cross-correlation output that falls below the threshold at the sample time.

requirements can be problematic. As noted previously, the center-to-center spacing between gratings must equal 1 cm for operation at 10 Gb/s. Due to their complexity, it can be difficult to manufacture sampled FBGs shorter than 1 cm that have the high reflectivity required to produce good correlation results. This problem can be resolved by inserting a passive splitter after the circulator and interleaving the gratings between multiple fiber branches. This decreases the spacing requirement by a factor equal to the number of branches. While these limitations on sampled FBG systems reduce the scalability of the architecture, a recent report details a new sampled FBG structure that can reduce the length of sampled FBGs while maintaining high reflectivity, enabling the application of this correlation technique to higher bit-rate systems [11].

To demonstrate multiwavelength correlation, we constructed an interleaved sampled FBG correlator using two seven-channel sampled gratings with 100-GHz channel spacing and a maximum reflectivity of 93%. As shown in Fig. 12, the sampled grating reflection spectrum has multiple passbands, one for each wavelength channel, that all shift together when the grating is stretched. Multiwavelength correlator operation is illustrated in Fig. 13, where the correlator is configured to match a "1x1" pattern. The first and third sampled FBGs are tuned to reflect and the second is tuned to zero reflectivity. At the correlator output, the signal is demultiplexed and the output correlation sequences for each wavelength are displayed on an oscilloscope. The 3-bit pattern on λ_1 is a match to the correlator and produces the expected "10 201" autocorrelation response, whereas the "100" pattern on λ_2 is not a match and produces a cross-correlation sequence that falls below the threshold at the sample time.

The experimental setup used to demonstrate multiwavelength correlation at 10 Gb/s is shown in Fig. 14. Four 53-byte pseudo-random NRZ packets with different 4-bit headers were modulated onto two wavelength channels (1555.6 and 1557.2 nm). Packet synchronization was assumed and a packet-rate clock signal was supplied to the decision circuits to sample the correlation peaks at the appropriate time. The correlators were first configured to match a header pattern of "1010." In the ones correlator, the first and third gratings are tuned via stretching to partially reflect, and the second and fourth gratings are tuned for no reflection. The gratings are configured in complement in the zeros correlator. Packets incoming on the two WDM channels are correlated by the gratings and sent to individual packet-rate decision circuits. The resulting match/no-match signals for each channel are used to control two $LiNbO_3$ optical switches, where header-matched packets are switched to port 1 and nonmatching packets are switched to port 2. The experimental results are shown in Fig. 15. The packets containing "1010" headers produce match signals that flip the switch and route these packets

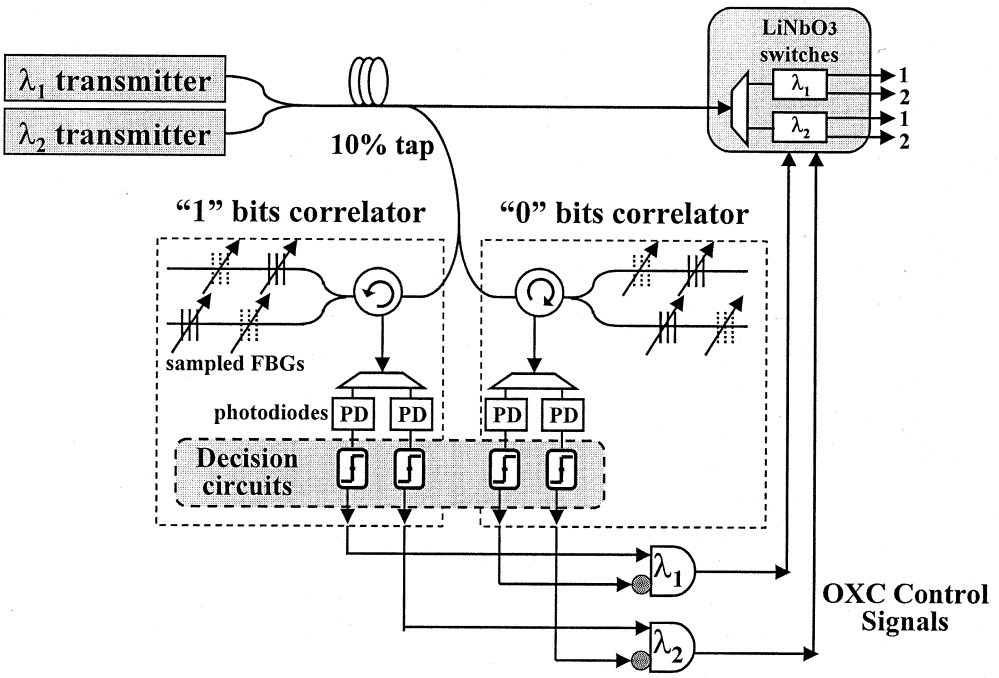

Fig. 14. Experimental setup for multiwavelength optical correlation using sampled FBGs. The decision outputs for the two wavelength channels are used as control signals for two optical switches. Packets with headers that produce a "match" signal are switched to port 1 and "no-match" packets are switched to port 2.

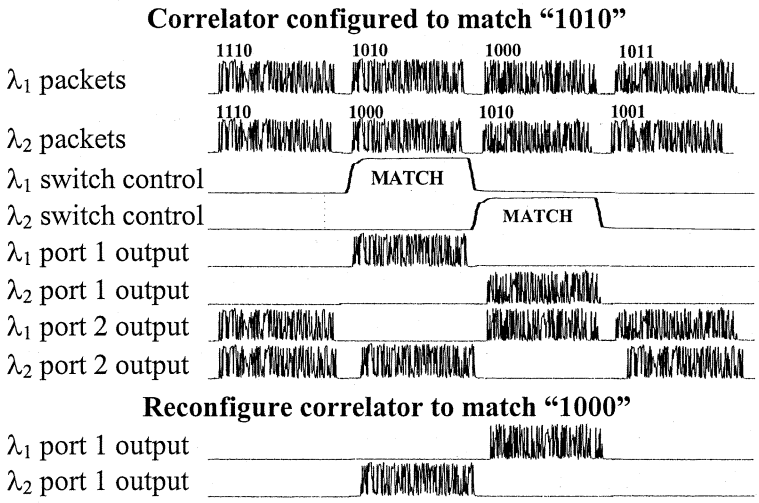

Fig. 15. Multiwavelength header correlation and switching results. Four packets incoming on two WDM channels are correlated with the pattern "1010." Matching packets are switched to port 1 of the optical switches and nonmatching packets are switched to port 2. The last two rows show the results when the correlators are reconfigured to match a "1000" pattern.

to port 1. The correlators are then reconfigured to recognize a "1000" pattern, causing the packets with these headers to now be routed to port 1. The total optical loss for the data through-path (90% tap + AWG demux + optical switch) was 7.7 dB, and the throughput data suffered zero power penalty when compared to the back-to-back receiver sensitivity at 10^{-9} bit error rate, as shown in Fig. 16.

A significant issue that arises for this interleaved correlator structure is the power fluctuations in the correlation output pulses due to coherent interference. The coherence time of standard telecommunication lasers is typically tens of nanoseconds, corresponding to a coherence length in fiber of about 2 m. When the differential time delay between two branches of the correlator is less than the coherence time of the laser (which is the case here), the recombined signals will coherently interfere with each other causing large power fluctuations in the correlation output function. This effect severely limits the stability of the correlation output and must be mitigated or prevented in order to effectively operate the correlator. There are a number of ways to resolve this problem. A polarization controller followed by a polarization beam splitter can be used at the input of each 1×2 optical splitter to ensure that the polarizations between the two branches are orthogonal. This will prevent coherent interference of the recombined

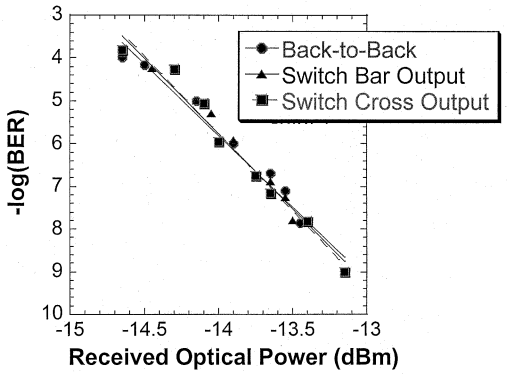

Fig. 16. Bit error rate (BER) performance of the throughput data packets. The circles indicate the baseline back-to-back performance (no switch) and the triangles and squares show the BER after passing through the "bar" and "cross" paths of the $LiNbO_3$ optical switch, respectively.

signals because orthogonally polarized light beams will not interfere. For more than two branches, a tree structure of 1×2 splitters with polarization controllers and polarization beam splitters can be used. Another, more manageable, solution is to somehow convert the coherent light into an incoherent signal before it enters the correlator. One method of doing this uses cross-gain modulation in a semiconductor optical amplifier (SOA) to transfer the coherent data pattern onto incoherent light, which for this case is the amplified spontaneous emission light generated by the SOA [12]. The single fiber correlators presented in the previous section have the advantage that we did not observe these coherent interference effects.

VI. Discussion

The use of thermally controlled FBGs as tunable-reflectivity mirrors to construct an optical correlator raises several important issues that must be considered, especially when scaling the design to operate at higher bit rates. Significant issues associated with device fabrication include polarization dependence, time-delay variation with reflectivity, grating bandwidth, dispersion, and thermal crosstalk between microheaters. These issues are simple to control at 10 Gb/s, scaling to 40 Gb/s appears quite feasible, and, with careful design, it should be possible to enable operation up to 100 Gb/s and perhaps beyond.

When writing the FBG onto which the heaters will be deposited, care must be taken to avoid any polarization dependence of the reflection spectrum; otherwise the reflectivity as well as the time delay for the grating will be different for the two orthogonal polarization modes of the optical signal. This can be avoided by carefully writing the grating so there are no asymmetries in the index-modulation profile. The 10-cm-long FBG that was fabricated for our experimental demonstration exhibited some polarization dependence, which made it difficult to adjust the subgrating reflectivities to produce the desired correlator configuration. A polarization controller was placed at the correlator input to optimize the levels of the output pulses while the correlator was being configured to recognize a desired bit sequence. However, because each subgrating is tuned to a different reflectivity, they each exhibit different polarization dependencies. For a particular tuning condition, the level and time delay of a pulse reflected off one subgrating will vary with polarization in a different way than the pulses from other subgratings. It is therefore challenging to locate a polarization that is simultaneously optimum for all the subgratings. Although this problem was not severe enough to prevent successful operation of the correlator, it did cause some distortion of the correlation outputs and should ideally be avoided when the grating is fabricated.

Using the edge of the FBG filter spectrum to adjust the reflectivity of the input signal causes dispersion and unwanted variations in the time delay as the grating is tuned. The round-trip delay of a signal reflected from an FBG varies for different points along the filter edge. This is a problem for our optical correlator, for which the time delay between neighboring gratings should ideally equal one-half of a bit time to ensure the output bits are exactly aligned to stack up and produce well-defined correlation peaks. In our experiment, we observed that the time delay of a pulse reflecting off a single subgrating varied as much as 35 ps, a significant fraction of the 100-ps bit time. Clearly, the time delay cannot vary more than the finite length of the grating (in the worst case, the light has shifted from effectively reflecting off the front of the grating to the back). The lengths of the subgratings are defined by the temperature profile induced by the heaters. Given that the length of our heaters is 0.5 cm and we expect that the heat diffuses approximately 1 mm beyond the heater edges (see thermal crosstalk discussion below), the subgrating lengths are $\sim$7 mm each. This corresponds to $\sim$34 ps in fiber, matching our measured results. At higher bit rates, the length of the heaters becomes much shorter, reducing the maximum possible delay variation, but this is in proportion to the shorter pulse-widths, and the variation will still be on the order of 25–35% of a bit time. This time-delay variation causes a spreading of the output autocorrelation peak and reduces the contrast between the central peak and the sidelobes, which adversely affects the threshold detection performance. Moreover, the time-delay response at the edge of an FBG passband is strongly sloped, causing the signal to experience dispersion. This was not a significant issue for our 10-Gb/s experiment, but it could become a problem at 40 Gb/s and higher. Note that no dispersion is induced on the throughput data channel, just on the tapped-off signal traversing the correlators. This causes the peaks in the correlation output to spread in time. Some amount of dispersion-induced spreading may actually be helpful for relaxing the required rise-time of the decision circuit's sampling signal. Though the decision circuit only needs to sample at the packet rate (MHz), at 40 Gb/s it must be able to recognize a correlation peak that is only 25 ps wide. Thus, it may be desirable to have some dispersion to relax this requirement; however, this must be carefully traded against the corresponding impairment to the contrast between the central peak and its sidelobes.

Another issue to consider is that the gratings must be strong enough to provide high reflectivity for the short subgratings as well as sufficient filter bandwidth for the incoming signal. The strength of a grating refers to the modulation depth δn of the variation in the fiber's effective index of refraction n_{eff}. Strong gratings with δn on the order of 10^{-3} are producible and should

meet our requirements for reflectivity. For strong gratings, the FBG bandwidth is directly proportional to δn and is independent of the grating length [13]. This must be considered at high bit rates to ensure that the filter bandwidth exceeds the bandwidth of the optical signal. For instance, 100-Gb/s pulses have an optical bandwidth of approximately 200 GHz. An FBG with $n_{\text{eff}} = 1.45$, a center wavelength of 1561 nm, and $\delta n = 2 \times 10^{-3}$ will yield a filter bandwidth of approximately 270 GHz, which is producible and wide enough to accommodate the high-bandwidth signal. However, stronger gratings have sharper filter edges, resulting in higher dispersion. Therefore, a careful design will have to trade these two parameters against each other to optimize the correlator performance at the desired bit rate.

To scale the correlator to higher bit rates requires a much higher spatial frequency of thin-film microheaters. However, thermal diffusion between heaters will limit how closely they can be placed before the thermal profile in the fiber between neighboring heaters begins to overlap, making it impossible to create a series of distinct subgratings. Researchers in [15] have investigated the possibility of using thin-film heaters to produce temperature distributions for the formation of superstructure long period gratings, which require a spatial modulation on the order of 0.5 mm. Their results showed that although it was possible to generate large gradients in temperature, the depth of modulation dramatically decreases as the spatial modulation period approaches 1 mm. This indicates that the heat induced by a thin-film heater on a standard fiber extends about 0.5 mm beyond its edges. However, this result does not explicitly consider the possibility of using etched cladding fibers to reduce the fiber diameter, which will strongly diminish the thermal diffusion and should enable the generation of periodic temperature variations with much higher thermal gradients [14]. For example, if we instead use a 30-μm-diameter etched cladding fiber, the 0.5-mm diffusion length is reduced to about 0.13 mm [14]. It should therefore be feasible to produce FBG correlators using heater arrays with spacings of a few hundred micrometers. Using etched cladding fibers, and with the possible addition of heatsinks in the spaces between heaters [14], [15], we expect that a 40-Gb/s correlator, which requires 1.25-mm spacings, should be readily producible, and a 100-Gb/s correlator, which requires 500-μm spacings, appears to be quite feasible.

An additional issue to consider when using optical correlators for header recognition is that the correlation outputs do not have reduced sidelobes like they would in coded communication schemes, such as optical CDMA (OCDMA). In OCDMA systems, the autocorrelation outputs look like a single tall peak with negligible sidelobes. This is because the codewords transmitted in OCDMA networks are designed to be orthogonal. This is necessary because the receivers in OCDMA networks are simultaneously receiving data packets from all of the users in the network, so orthogonal codes must be employed to ensure that the cross-correlation outputs from other users do not add up and exceed the threshold at the sample time. For header recognition, we only receive one bit pattern at a time and there is no option for demanding that Internet addresses use orthogonal codes. Thus, correlation output functions in header-recognition systems will contain sidelobes. An important consideration for these systems is the signal-to-noise ratio so that a sidelobe from a nonmatching packet does not inadvertently exceed the threshold at the sample time. A valuable next step for this research topic would be to perform a full analysis of the correlation error rates (how often the correlator produces either false positive or negative matches) that result for different threshold levels and for different random input addresses.

VII. Conclusion

Fiber-based optical correlation techniques have been investigated for several years for their potential to recognize incoming bit streams at the speed of light, with essentially no latency. Present optical technologies are limited to testing an incoming bitstream against a handful of different possible correlation sequences with only a few bits each. This limits their potential uses for header recognition since real network data packets have addresses with dozens of bits that must be compared against thousands of entries in a routing lookup table. However, our research indicates that it may be feasible to design an algorithm to generate a subset of the routing table containing entries for packets that can be forwarded by examining only a few bits in an IP destination address. This opens the door to using the best that both optics and electronics has to offer and enables the construction of an optically assisted Internet router.

In this paper, we presented the constraints and optimization goals for an algorithm responsible for generating a reduced optical lookup table that aims to maintain a high hit rate while minimizing the number of correlators and bits/correlator that are required. To implement the optical bypass, we presented a novel FBG-based correlator design in which an array of thin-film microheaters is deposited on a long, uniform FBG to create a series of equally spaced subgratings that are tuned by varying the voltages across the heaters. The advantages of this design are that the fiber correlators are compact, electrically tunable, simple to produce, and readily scalable to higher bit rates. An 8-bit correlator module was constructed and used to experimentally demonstrate the successful correlation and switching of packets at 10 Gb/s. With this technique, the packets that successfully find a match in the bank of optical correlators can be forwarded at nanosecond speeds, an order of magnitude reduction from electronic lookup times.

To reduce the number of optical correlators that would be required to employ this optical bypass in a WDM routing node, we also constructed an optical correlator using sampled FBGs to enable the packet headers on different input wavelengths to simultaneously be compared with a correlation pattern.

References

[1] P. Mehrotra and P. D. Franzon, "Novel hardware architecture for fast address lookups," *IEEE Commun. Mag.*, pp. 66–71, Nov. 2002.

[2] J. Bannister, J. Touch, P. Kamath, and A. Patel, "An optical booster for internet routers," in *Proc. 8th Int. Conf. High Performance Computing*, Hyderabad, India, Dec. 2001, pp. 339–413.

[3] J.-D. Shin, M.-Y. Jeon, and C.-S. Kang, "Fiber-optic matched filters with metal films deposited on fiber delay-line ends for optical packet address detection," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 941–943, July 1996.

[4] D. B. Hunter and R. A. Minasian, "Programmable high-speed optical code recognition using fiber Bragg grating arrays," *Electron. Lett.*, vol. 35, pp. 412–414, Mar. 1999.

[5] J. E. McGeehan, S. Kumar, J. Bannister, J. Touch, and A. E. Willner, "Optical time-to-live decrementing and subsequent dropping of an optical packet," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2003, FS6.
[6] A. Othonos and K. Kalli, *Fiber Bragg Gratings: Fundamentals and Applications in Telecommunications and Sensing*. Boston, MA: Artech House, 1999, p. 99.
[7] R. Kashyap, *Fiber Bragg Gratings*. San Diego, CA: Academic, 1999, pp. 436–440.
[8] M. C. Cardakli and A. E. Willner, "Synchronization of a network element for optical packet switching using optical correlators and wavelength shifting," *IEEE Photon. Technol. Lett.*, vol. 14, pp. 1375–1377, Sept. 2002.
[9] L. Li, J. Geng, L. Zhao, G. Chen, G. Chen, Z. Fang, and C. F. Lam, "Response characteristics of thin-film-heated tunable fiber Bragg gratings," *IEEE Photon. Technol. Lett.*, vol. 15, pp. 545–547, Apr. 2003.
[10] M. Ibsen, M. K. Durkin, M. J. Cole, and R. I. Laming, "Sinc-sampled fiber Bragg gratings for identical multiple wavelength operation," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 842–844, June 1998.
[11] N. Yusuki and Y. Shinji, "Realization of various superstructure fiber Bragg gratings for DWDM systems using multiple-phase-shift technique," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2002, TuQ3.
[12] P. Parolari, L. Marazzi, D. Rossetti, G. Maier, and M. Martinelli, "Coherent-to-incoherent light conversion for optical correlators," *J. Lightwave Technol.*, vol. 18, pp. 1284–1288, Sept. 2000.
[13] T. Erdogan, "Fiber grating spectra," *J. Lightwave Technol.*, vol. 15, pp. 1277–1294, Aug. 1997.
[14] E. R. Lyons, "Tunable all-fiber devices for optical fiber communications," Ph.D. dissertation, Univ. of California, Irvine, 2001.
[15] J. A. Rogers, P. Kuo, A. Ahuja, B. J. Eggleton, and R. J. Jackman, "Characteristics of heat flow in optical fiber devices that use integrated thin-film heaters," *Appl. Opt.*, vol. 39, pp. 5109–5116, Oct. 2000.
[16] G. G. Finn, S. Hotz, and C. Rogers, "Method and networking interface logic for providing embedded checksums," U.S. Patent 5 826 032, Oct. 20, 1998.

Michelle C. Hauer (S'96) received the B.S. degree in engineering physics from Loyola Marymount University, Los Angeles, CA, in 1997 and the M.S.E.E. degree from the University of Southern California (USC), Los Angeles, in 2000, where she is currently pursing the doctoral degree.

She is currently a Research Assistant in the Optical Fiber Communications Laboratory, USC. She is also a Systems Engineer for Raytheon Space and Airborne Systems, El Segundo, California. Her research interests include optical signal-processing techniques for implementing all-optical networking functions, polarization effects in lightwave systems, and quantum communications. She is a contributing author on more than 20 research papers.

Ms. Hauer is a Member of Tau Beta Pi, Eta Kappa Nu, and Sigma Pi Sigma. She is a Student Member of IEEE LEOS, the Optical Society of America (OSA), and SPS.

John E. McGeehan (S'96) received the B.S. and M.S. degrees in electrical engineering from the University of Southern California (USC), Los Angeles, in 1998 and 2001, respectively, where he is currently pursuing the Ph.D. degree.

He joined the Optical Communications Laboratory, USC, as a Research Assistant in 2001. His research interests include the implementation of all-optical networking functions and optical signal processing as well as Raman amplification and signal monitoring. He is the author or coauthor of 12 technical papers

Saurabh Kumar (S'03) received the B.E. degree in electrical and electronics engineering from the Birla Institute of Technology and Science, Pilani, India, in 2001. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Southern California (USC), Los Angeles.

He is a Research Assistant in the Optical Communications Laboratory, USC. His research focuses on all-optical signal processing to enable the implementation of networking functions in the optical domain.

Mr. Kumar is a Member of Eta Kappa Nu and a Student Member of OSA.

Joseph D. Touch (S'84–M'91–SM'02) received the B.S. degree (Hons.) in biophysics and computer science from the University of Scranton, Scranton, PA, in 1985, the M.S. degree from Cornell University, Ithaca, NY, in 1987, and the Ph.D. degree from the University of Pennsylvania, Philadelphia, in 1992, both in computer science.

He joined the USC/Information Sciences Institute, Marina del Rey, CA, in 1992 and is currently Director of the Postel Center for Experimental Networking in the Computer Networks Division. He is currently involved with fault-tolerant networks such as DynaBone, NetFS, and X-Tend), optical Internet wide-area networks (WANs) and local-area networks (LANs) such as POW and OCDMA, geographic routing such as GeoNet, and smart space devices for user presence. He has led projects ranging from gigabit LANs (such as ATOMIC2), NIC design (such as PC-ATOMIC), multicast web caching (I.SAM), to his most recent past project in the automated deployment and management of virtual networks (XBONE). He is also a Research Associate Professor in the Department of Computer Science, University of Southern California, where he taught advanced operating systems for several years. He now runs a program for summer students (SGREP) and advises a number of graduate students. He has several patents current and pending and is the author of a number of publications, including the book *High Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication*. His primary research interests include virtual networks, automated network management, high-speed protocols, empirical protocol performance analysis, Internet architecture, and protocols for latency reduction.

Dr. Touch is a Member of Sigma Xi (A'84–M'93) and ACM (S'83–M'92), Co-Chair of the IEEE ITC (Internet) committee, and is active in the IETF. He also serves on the editorial boards of IEEE NETWORK and Elsevier's *Computer Networks*, and is a Member of several program committees, including IEEE INFOCOM (since 1994) and Sigcomm, and was Vice General Chair of Opticomm 2001.

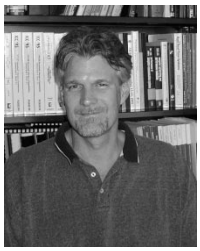

Joseph Bannister (S'80–M'80–SM'95) received the B.A. degree (with high distinction) in mathematics from the University of Virginia, Richmond, in 1977 and the M.S. degree in electrical engineering and the M.S. and Ph.D. degrees in computer science from the University of California, Los Angeles, in 1980, 1984, and 1989, respectively.

He previously held positions at The Aerospace Corporation, System Development Corporation (now Unisys), Sytek (now Hughes Network Systems), Research Triangle Institute, and Xerox. He is currently a Director of the Computer Networks Division, Information Sciences Institute, and a Research Associate Professor in the Department of Electrical Engineering-Systems for the University of Southern California in Los Angeles. His current technical interests are in high-speed and optical networking, routing, network management and design, performance evaluation, and protocol engineering. He has more than 50 publications in high-speed networking, distributed computing, and network management. He serves on the editorial boards of *Optical Networks* and *Computer Networks*.

Dr. Bannister is a Member of Sigma Xi, the Association for Computing Machinery (ACM) SIGCOMM, the Internet Society, and the American Association for the Advancement of Science (AAAS). He actively participates in the research community, serving or having served on the program committees or as Chair of INFOCOM, Interop, ICNP, ICCCN, LAN/MAN Workshop, OPTICOMM, and the IEEE/TCCC Computer Communications Workshop. He is or has been the Principal Investigator of the Defense Advanced Research Projects Agency (DARPA) and National Science Foundation (NSF) projects on decentralized network management, metacomputing, optical networks, quality of service management, wireless communications, satellite communications, battlefield awareness and data dissemination, lambda grids, and network vulnerability analysis. In 1996, he participated on a DARPA ISAT Study Group on Network Survivability, which resulted in the establishment of a major DARPA program in information assurance and survivability. In 2000 and 2003, he served on the NSF CISE ANIR Committee of Visitors, which was responsible for reviewing the foundation's research program in networking. He is currently serving a term on the NSF CISE Advisory Committee, advising the CISE Assistant Director on matters related to the division's research portfolio. He also co-chaired a panel on developing DARPA's long-term networking and communications research and development strategy. He has supported NSF, DARPA, and DOE on source selections, review panels, and other advisory tasks.

Edward R. Lyons (M'84) received the B.S. degree from the University of California, Los Angeles, in 1983 and the M.S. and Ph.D. degrees from the University of California at Irvine in 1994 and 2001, respectively.

He has more than 15 years of experience at Raytheon Space and Airborne Systems and Hughes Aircraft Company working on electrooptical and photonic systems. Since 1987, he has been involved in photonics systems development at Hughes and Raytheon sites throughout southern California. Photonic projects include work on secure fiber-optic communications systems, packaging of $LiNbO_3$ modulators, and numerous programs involving high-performance analog fiber-optic link development. His research interests include tunable all-fiber devices and quantum communications.

C. H. Lin, photograph and biography not available at the time of publication.

A. A. Au, photograph and biography not available at the time of publication.

H. P. Lee, photograph and biography not available at the time of publication.

Dmitry S. Starodubov (M'03) received the M.Sc. degree in optical engineering from the Moscow Institute of Radioengineering, Electronics and Automation, Russia, in 1992 and the Ph.D. degree in laser physics from the General Physics Institute, Russian Academy of Sciences, in 1995.

He received a grant from the National Research Council and was a CAST Fellow at the University of Southern California from 1996 to 1997. In 1998, he cofounded Sabeus (formerly D-STAR Technologies), where he presently serves as the Chief Technology Officer. His research interests include the physics of glasses, nonlinear optics, optical communications, and fiber sensors. He is an author of 15 patents and more than 100 publications.

Alan Eli Willner (S'87–M'88–SM'93–F'04) received the B.A. degree from Yeshiva University, New York, and the Ph.D. degree in electrical engineering from Columbia University, New York.

He has been with AT&T Bell Labs and Bellcore, and is a Professor of electrical engineering at the University of Southern California. He was Photonics Division Chair of OSA; Program Cochair of the OSA Annual Meeting; Program Cochair of CLEO; Steering and Technical Committee of OFC, Program committee member of ECOC. He has more than 300 publications, including one book. His research is in optical fiber communication systems.

Prof. Willner is a Fellow of the Optical Society of America (OSA) and was a Fellow of the Semiconductor Research Corporation. He has received the NSF Presidential Faculty Fellows Award from the White House; the David and Lucile Packard Foundation Fellowship; the NSF National Young Investigator Award; the Fulbright Foundation Senior Scholar Award; the IEEE LEOS Distinguished Lecturer Award; the USC/Northrop Outstanding Junior Engineering Faculty Research Award; the USC/TRW Best Engineering Teacher Award; and the Armstrong Foundation Memorial Prize. He was Vice-President for Technical Affairs for IEEE LEOS; a member of the LEOS Board of Governors; Cochair of the OSA Science and Engineering Council; General Chair of the IEEE LEOS Annual Meeting; Editor-in-Chief of the JOURNAL OF LIGHTWAVE TECHNOLOGY; Editor-in-Chief of the IEEE JOURNAL OF SELECTED TOPICS IN QUANTUM ELECTRONICS; Guest Editor of the JOURNAL OF LIGHTWAVE TECHNOLOGY Special Issue on wavelength-division multiplexing; and Guest Editor of the IEEE JOURNAL OF QUANTUM ELECTRONICS.