Dynamically Reconfigurable All-Optical Correlators to Support Ultra-fast Internet Routing

M. C. Hauer, J. McGeehan, J. Touch*, P. Kamath*, J. Bannister*, E.R. Lyons+, C. H. Lin+, A. A. Au+, H.P. Lee+, D. S. Starodubov, and A. E. Willner**

Dept. of Electrical Engineering-Systems, University of Southern California, Los Angeles, CA 90089-2565
(213) 740-1488, Fax (213) 740-8729, hauer@usc.edu
**USC Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA 90292-6695*
+Dept. of Electrical and Computer Engineering, UC Irvine, 2231 Engineering Gateway, Irvine, CA 92697
***Sabeus Photonics, 3995 Via Oro Avenue, Long Beach, CA 90810*

Abstract: A new all-optical networking function is demonstrated using a bank of novel thermally-tuned FBG-based optical correlators to construct an "optical bypass" to vastly accelerate conventional electronic internet routers. Based on a simple set of rules, a software algorithm configures the correlators as a routing table cache that can quickly determine the destination port for a large percentage of the incoming traffic by only examining a small subset of the address bits.

OCIS codes: (060.2330) Fiber optics communications; (060.4250) Networks

Introduction

As transmission speeds in the network core approach 40Gb/s and beyond, the need to make faster routing decisions at each networking node becomes increasingly evident. It is therefore important for optical designers to examine how current electronic routers function and determine where it may be feasible to employ optical techniques to aid the electronics in making ultra-fast routing decisions.

In conventional Internet routers, packets are steered towards their destinations by interrogating their 24-bit destination IP addresses and matching them to entries in a large routing table using complicated "longest prefix" matching algorithms. This can be time consuming given that core routing tables contain upwards of 100,000 entries. As a result, a true all-optical router would need to be capable of 24-bit lookups into 100,000-entry tables at ≥40 Gb/s. Such capabilities are beyond current optical technologies, but some recent developments hint at the feasibility of a partial solution. Given that most core routers have only four to eight outgoing ports, it may be possible to determine a packet's outgoing 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 vastly 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. The optical bypass can utilize a subset of the routing table with as few as 100 entries and still successfully optically route as much as 90% of the incoming traffic [1]. The remaining challenge is to determine how to design a 24-bit input, 100-entry optical index using a manageable number of optical correlators.

The optical correlators required for this application must be tunable and designed to easily scale to 40Gb/s and beyond. Optical correlators operate by sequentially splitting the optical bit pattern and recombining it after one branch has experienced a one to few bits time delay. The tiny distances corresponding to these delays would present a serious challenge for previously reported correlators which, for example, are implemented with either optical splitters and fiber mirrors [2] , or an array of discrete fiber Bragg gratings (FBG) tuned with separate piezoelectric stretchers [3]. In these cases, the inter-mirror or inter-grating spacings would have to be 2.5 mm for a one-bit round trip time delay at 40Gb/s – an impractical length for a device using discrete fiber components. Thus, we propose a new 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 micro-heaters. Millimeter, and even micron sized spacings are easily achieved with this technology. An additional advance over previously reported correlators is that our approach uses two grating arrays per correlator, one to correlate with the "1-bits" in the desired code and the other for the "0-bits," allowing all possible input codes to produce unique correlation outputs.

Concept

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 fiber-Bragg-grating based optical correlators. A conceptual

diagram showing how the optical bypass is implemented in an IP router is shown in Fig. 1. 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 into the multiple correlators, each of which can be configured to produce a "match" signal for any number, X ,of 24 bits. For example, if the algorithm has predetermined that any incoming packets with bit positions 1, 4, and 5 equal to 1, 0, and 1 should go to port 1, then the correlator is configured to provide a "match" signal for any input pattern with "1xx01" for its first five bits and anything for the rest (where the "x"s indicate a "don't care" bit). There will be a group of correlators for each output port since there will be a group of bit patterns that the software has determined should all be routed to a particular port. The goal of the algorithm is to reduce the number of correlators required. Threshold detectors are used at the outputs of the optical correlators to provide an electrical match/no match condition to the optical switch. The switch uses these signals to determine which output port each packet should be routed to. For any packet that the correlators fail to find a match for, the switch sends it to an auxiliary port for electrical processing by a conventional electronic router.

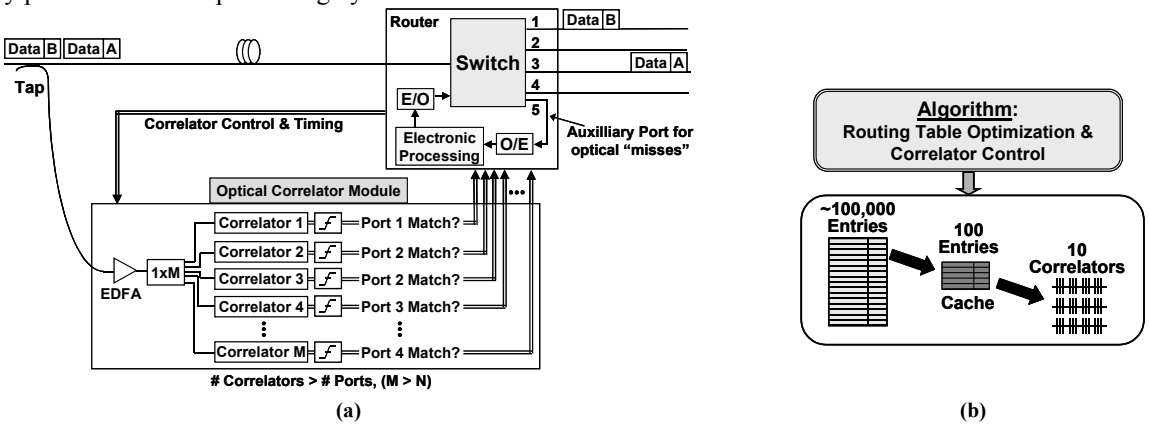

Figure 1. (a) Conceptual diagram of an internet core router with **N** output ports (typically 2 to 4) assisted by a bank of **M** optical FBG-based correlators that are dynamically configured by a software algorithm. (b) The algorithm first reduces the large routing table down to a cache of the 100 most popular entries. Then, after grouping these entries by output port, each correlator is configured to match to a particular group by recognizing only a subset, **k**, of the **n** bits in the destination address. If the correlator fails to match the incoming packet address to a group, then the packet is processed by conventional electronics.

Routing Algorithm

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 $F:\{0,1\}^n \to \{1,2,\ldots,N\}$ denote the routing table, which is a function that maps any destination address to its output interface. The *i*th 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 $1 \le i \le M$ and $1 \le j \le n$. 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; hence $|B_i| \ge 2^{n-K}$ for $1 \le i \le M$. 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: $|B_i| \ge 2^{n-K}$
3. Nonoverlapping correlators: if $i \ne j$ then $B_i \cap B_j = \varnothing$

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 $N = 8$ (the number of output ports in a typical Internet router), $M = 8$, $n = 24$ [the most-significant bits of an IP v4 (Internet Protocol Version 4) 32-bit address], and $K = 5$ (a constraint of the technology). The major contributions of this paper are that it is feasible to implement a bank of optical correlators that will meet the requirements above.

Experimental Setup

The experimental setup to demonstrate the optical bypass at 10Gb/s with 53 byte NRZ packets is shown in Fig. 3(a). Although this design is best suited for higher bit rates, ≥40 Gb/s, the demonstration is performed at 10Gb/s due to equipment limitations. The optical correlators were fabricated using e-beam evaporation to deposit an array of 8 metallic heaters with a period of 0.5cm onto two 10cm long uniform fiber Bragg gratings at a center wavelength of ~1561nm. The heaters are composed of 15nm of titanium and 120nm of gold. The gratings were fabricated with Δn

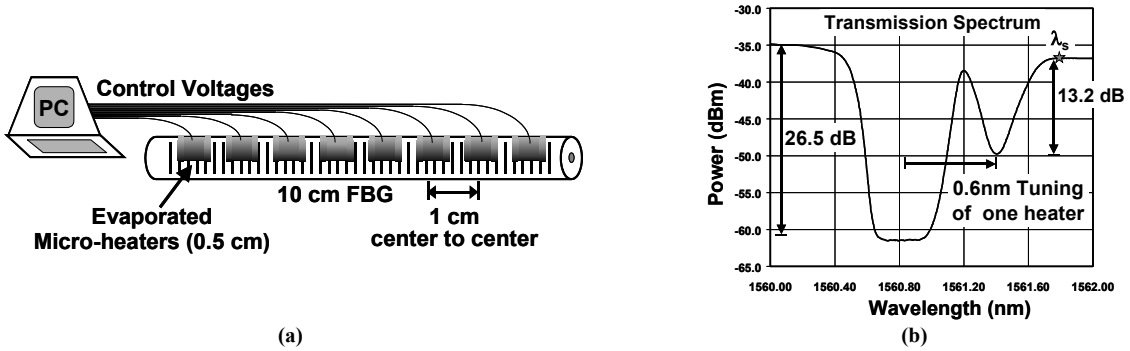

Figure 2. (a) Fiber Bragg grating correlator – 8 evaporated micro-heaters, each 0.5 cm in length with 0.5 cm spaces between them, deposited onto one 10 cm long, strong, uniform FBG. (b) Optical spectrum of FBG in transmission showing one heater being tuned 0.6 nm away from the main, 10 cm FBG, peak. For the experiment, the signal wavelength was 1561.80 nm which required 1 nm tuning to put the signal at max reflectivity (~95% reflectivity) which required ~ 25 mW tuning power.

~ 10^{-3} to provide >90% reflectivity when only 0.5cm sections were used. The center wavelength of the gratings is chosen to be out-of-band so that only those sections that correspond to bits that need to reflect are tuned into band. To configure the correlator to recognize a bit pattern of, say, 0x1x1xxx (only 3 of 8 bits matter), the third grating in the "1 bits" correlator is tuned to partially reflect and the fifth grating is tuned for full reflection and the first grating in the "0 bits" correlator is also tuned for full reflection. Thermal tuning is accomplished by applying a voltage across the micro-heater. The threshold range of the "1 bits" decision circuit is set to detect two stacked bits whereas the "0 bits" threshold is always set to detect a 0 level. A packet-rate timing signal is used to trigger the decision circuits to sample the correlator outputs at the right time. In a full system implementation, this timing signal would be generated by a previous module that detects the packet arrival time [4]. The two outputs are sent through an AND gate that provides a high signal when there is a match and a low signal otherwise. This output can be used to configure a switch to route the packet to the correct port.

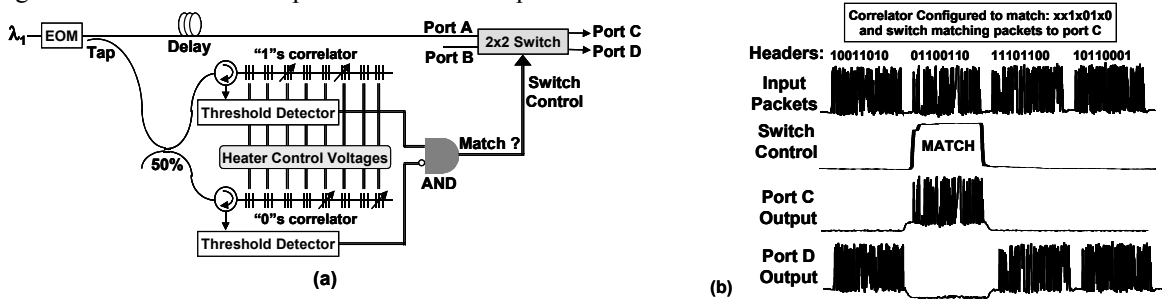

Figure 3. (a) Experimental setup of an 8-bit optical correlator tuned to recognize an xx1x01x0 bit pattern, and (b) routing results showing the successful recognition and switching incoming packets to the outputs of a 2-port switch.

Results

To demonstrate the successful operation of the optical bypass function, the FBG correlators were tuned to recognize an xx1x01x0 (4 of 8 bits) pattern and the correlation outputs were used to route packets with matching headers to Port C of a two port $LiNbO_3$ optical switch. Fig. 3(b) depicts the input packet stream, the switch control signal from the correlator output, and the final switching results, showing the successful routing of the packets.

References

[1]J. Bannister, J. Touch, P. Kamath, A. Patel, "An Optical Booster for Internet Routers," Inviited Paper, Proc. Eighth Int'l Conf. on High Performance Computing, to appear Dec. 2001

[2]J.-D. Shin et al, IEEE Photon. Tech. Lett., **8**, 941-943, (1996)

[3]D. B. Hunter et al, IEEE Electron. Lett., **35**, 412-414, (1999)

[4]M. Cardakli et al, Optical Fiber Comm. Conference, **2**, paper TuW2 (2001)