10.1117/2.1200809.1294

Components developed for all-optical Internet router

Joseph Touch

Internet routers pass packets through a series of specific functions that can be implemented more readily in optics using high-speed serial methods than by translating conventional parallel electronics.

Internet routers direct packets, formatted blocks of data, to their destination. Today's network core routers handle links up to 40 gigabits per second using expensive and power-hungry parallel electronics. Optical routers could support much higher rates more efficiently and simply if existing challenges could be overcome.[1] For example, current electronics are far too complex to implement directly in optics, and no optical random-access storage exists. Most deployed optical networks use circuits, and recent work has tended to focus on managing these circuits. Our approach, on the other hand, uses the native capabilities of optics to provide a viable path toward an all-optical router.

Today's optical networking is based on optical channels that allocate circuits on a per-wavelength basis. Optical packet systems direct trains of packets, that are switched onto a wavelength circuit as a group, to a single destination (e.g., optical burst switching[2]). Both circuits and packet bursts represent intermediate stages in the evolution of the technology towards true packet switching, the basis of the Internet, which is more robust and efficient. Our research attempts to leapfrog this evolution for native optical Internets, focusing on packet services from the start.

In a process called 'forwarding', routers direct packets through a series of functions and then hand them off to the next router along a path (see Figure 1). A packet is first matched against a set of policy filters, and its next-hop destination is determined by an address match. The packet's time to live (TTL) field is decremented, and zero-TTL packets are dropped so that loops do not clog the network. A header checksum is updated, using ones-complement arithmetic. Competing packets are queued and transmitted later.

Our team has translated the matching, decrementing, and checksum functions into native optics. In each case, we focused on ultra-fast serial operations using simple components,

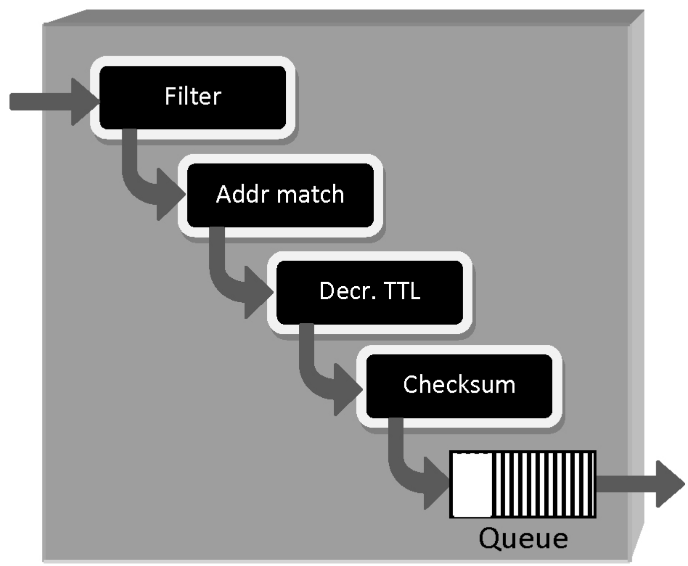

Figure 1. *A packet traverses a series of functions in an Internet router. Addr: Address. Decr.: Decrement. TTL: Time to live.*

rather than the conventional electronic approach of extreme parallelization. Our filter and address matching systems use pairs of complementary optical correlators that check for zeroes and ones in the pattern separately.[3] Our decrementer mechanism uses a single latch and comparator.[4] Our Internet checksum mechanism uses a single, two-input full adder. The result passes packets through only a handful of optical components rather than the myriad of parallel electronics currently required.

Consider decrementing the 8-bit TTL field. Electronic systems perform true arithmetic subtraction, using fast parallel hardware (see Figure 2a). Optics can support very fast serial operations, but only if they are very simple. We recognized that 'decrementing' is equivalent to inverting the bits of the field, least-significant bit first, until the first '1' is encountered.[4] This invert-and-hold mechanism can be implemented with an inverter, a switch, and a latch (see Figure 2b). We developed a similar technique to compute the ones-complement Internet checksum us-

Continued on next page

10.1117/2.1200809.1294 Page 2/2

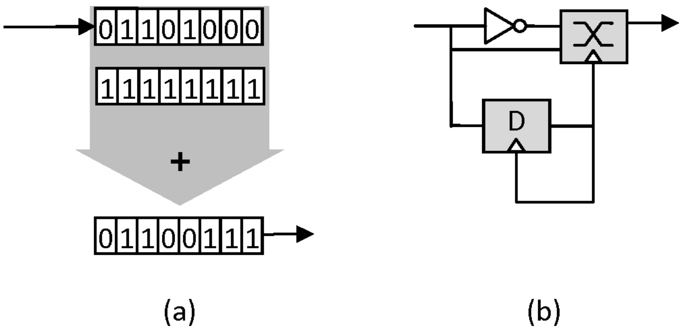

Figure 2. *Decrementing an 8-bit counter in electronics, using (a) a complex full 8-bit adder and (b) our serial optical solution.*

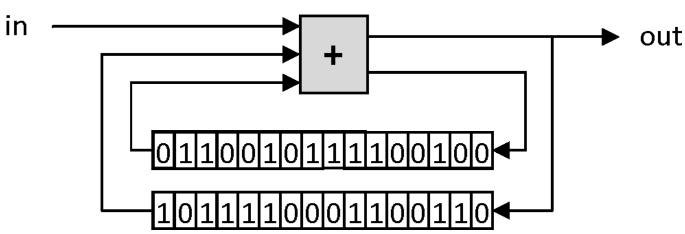

Figure 3. *Computing the Internet checksum in serial.*

ing recirculation buffers and a single full-bit adder (see Figure 3). These optical designs are infeasible in corresponding electronics at emerging line speeds.

We also avoid forwarding every possible packet. Current electronic routers have a 'fast path' for common packet formats and a general-purpose central processing unit (CPU) 'slow path' for rare, complex packets. Similarly, we apply optics to the majority of traffic, using a conventional electronic router as a bypass.[5] For example, we compute shared bit patterns for optical correlator matches, handling the majority of cases using a small number (8–10) of patterns, and passing unmatched packets to the bypass router. In comparison, handling every case of a 250,000-entry forwarding table using patterns would be intractable (e.g., 100 patterns).

The components we have developed demonstrate that true packet routing is not beyond the reach of existing optical technology. The two remaining challenges are integration and compensation for the lack of true optical queuing. Our devices use individual components, and single-chip integration is not yet feasible. Advances in filters, using tunable microdisks rather than fiber Bragg gratings, may present a way forward. We are also reexamining the notion of queuing, to determine if another, more restricted reordering (that optics can support) may be sufficient.

This work was partially supported by the Defense Advanced Research Projects Agency's Optical Code Division Multiple Access (OCDMA) and Packet over Wavelengths (POW) projects, and the National Science Foundation (NSF) Optical Networks initiative.

Author Information

Joseph Touch
Postel Center
University of Southern California Information Sciences Institute
Marina del Rey, CA
Departments of Computer Science and Electrical Engineering
University of Southern California
Los Angeles, CA
http://www.isi.edu/touch
http://www.postel.org/

Joseph D. Touch is the director of the Postel Center in the Informatin Sciences Institute (ISI) at the University of Southern California (USC) and a research associate professor in USC's computer science and electrical engineering departments. He also assists with the United States Air Force's Transformational Communications Satellite (TSAT) satellite network program. His research interests include network architecture, virtual networks, optical Internets, transport protocols, and network performance.

References

1. H. J. Chao, *Next generation routers*, **Proc. IEEE 90** (9), 2002.
2. C. Qiao and M. Yoo, *Optical burst switching – A new paradigm for an optical Internet*, **J. High Speed Networks 8** (1), pp. 69–84, 1999.
3. M. Hauer, J. McGeehan, S. Kumar, J. Touch, J. Bannister, E. Lyons, C. and A. Lau, H. Lee, D. Starodubov, and A. Willner, *Optically-assisted Internet routing using arrays of novel dynamically reconfigurable FBG-based correlators*, **IEEE/OSA J. Lightwave Technol. 21** (11), pp. 2765–2778, 2003.
4. J. McGeehan, S. Kumar, D. Gurkan, S. Motaghian Nezam, J. Bannister, J. Touch, and A. Willner, *All-optical decrementing of a packet's time-to-live (TTL) field and subsequent dropping of a zero-TTL packet*, **IEEE/OSA J. Lightwave Technol. 21** (11), pp. 2746–2752, 2003.
5. J. Bannister, J. Touch, P. Kamath, and A. Patel, *An optical booster for Internet routers*, **Proc. 8th Int'l IEEE/ACM Conf. High Perform. Comput.**, pp. 339–413, 2001. Also U.S. Patent 7,369,766, Nov. 2002.