All-Optical Decrementing of a Packet's Time-to-Live (TTL) Field and Subsequent Dropping of a Zero-TTL Packet

J. E. McGeehan, *Student Member, IEEE*, Saurabh Kumar, *Student Member, IEEE, Student Member, OSA*,
Deniz Gurkan, *Student Member, IEEE, Student Member, OSA*,
S. M. R. Motaghian Nezam, *Student Member, IEEE, Student Member, OSA*,
Alan Eli Willner, *Fellow, IEEE, Fellow, OSA*, K. R. Parameswaran, *Member, IEEE, Member, OSA*,
M. M. Fejer, *Member, IEEE*, J. Bannister, *Senior Member, IEEE*, and Joseph D. Touch, *Senior Member, IEEE*

***Abstract*—We demonstrate an optical time-to-live (TTL) decrementing module for optical packet-switched networks. Our module acts on a standard NRZ-modulated binary TTL field within a 10 Gb/s packet and decrements it by one if the TTL is nonzero. If the TTL of the incoming packet is zero, the module signals an optical switch to drop the packet. Our technique is independent of the TTL length, does not require the use of ultrashort RZ optical pulses, requires no guard time between the end of the TTL field and the packet data, and has only a 2.4 dB power penalty at 10^{-9} bit-error rate.**

***Index Terms*—Optical communications, optical networking, optical TTL, time-to-live (TTL).**

I. INTRODUCTION

THERE has been increased recent interest in the vision of an optical packet-switched network. In these networks, packets are processed and switched optically, without need for conversion of the optical packets to electronic signals at intermediate nodes within the network core. However, while all-optical packet switching remains a laudable goal for efficient and high-throughput networking, a move to optical packet switching within the network core may require implementation of a number of packet-processing and network-management functions currently performed by electronic routers, in the optical domain – a fact frequently ignored when the concept of the "all-optical network" is discussed.

One problem faced by many packet-switched networks is that of "routing loops," where misdirected or mislabeled packets are routed in circles, never reaching their destination, and leading to severe network congestion [1]. While rare packet processing errors can result in misrouted packets, routing loops can occur due to errors within the routing table stored at a switching node, or because a change in the routing table is taking longer than normal to propagate throughout the network. One commonly-employed method to prevent these loops from strangling the bandwidth of a network is to use a "time-to-live" (TTL) field within the header of a packet. The TTL field determines the maximum numbers of hops a packet can take before getting dropped for being too "old" for the network. In modern IP packets, the TTL is a binary number (typically 8 b long) that is decremented by 1 when traversing a switching node. When the TTL value of a packet reaches zero, the packet is dropped from the network. "Rogue" packets are either eventually rerouted toward the correct destination, or the TTL expires and they are dropped. Decrementing this TTL field requires editing the packet header, something easy in electronics but often difficult in optical systems.

There has been minimal previous research on optical technologies for decrementing the TTL field of an optical packet, or of dropping or destroying packets that have a zero TTL, particularly for NRZ systems. There has been one recent report of using a discrete series of ultrashort RZ optical pulses as a packet's "TTL burst," where instead of a binary field within the header the number of optical pulses present corresponded to the value of the TTL [2]. The TTL module, consisting of a semiconductor-optical-amplifier (SOA) and a mode-locked laser, would extinguish a single pulse each time the packet passed through the module. When no pulses remained, the packet was dropped. However, this method required restructuring of a standard packet header as the binary TTL field was replaced with an indeterminate number of optical pulses.

We demonstrate a module that acts upon a standard 8-b TTL field located within an NRZ-modulated optical packet. If the TTL field within the incoming packet is nonzero, the packet passes through the module, where gain saturation within an SOA and difference-frequency-generation (DFG) in a set of periodically poled lithium-niobate (PPLN) waveguide wavelength shifters are used to decrement the TTL field. If the

---

Manuscript received December 17, 2002; revised July 25, 2003. This work was supported in part by the National Science Foundation and the Intel Corporation.

J. E. McGeehan and S. M. R. M. Nezam are with the Department of Electrical Engineering–Systems, University of Southern California, Los Angeles, CA USA 90089 (e-mail: mcgeehan@usc.edu).

S. Kumar is with the Optical Communications Laboratory, University of Southern California, Los Angeles, CA 90089 USA (e-mail: skumar@usc.edu).

D. Gurkan is with the Information Sciences Institute, University of Southern California, Marina Del Rey, CA 90292 USA, and also with California State University Long Beach, Long Beach, CA 90840 USA (e-mail: denizgurkan@ieee.org).

A. E. Willner is with the Optical Fiber Communications Laboratory, University of Southern California, Los Angeles, CA 90089 USA (e-mail: willner@usc.edu).

K. R. Parameswaran and M. M. Fejer are with the Edward L. Ginzton Laboratory, Stanford University, Stanford, CA 94305 USA.

J. Bannister and J. D. Touch are with the USC Information Sciences Institute, Marina Del Rey, CA 90292 USA.

Digital Object Identifier 10.1109/JLT.2003.819131

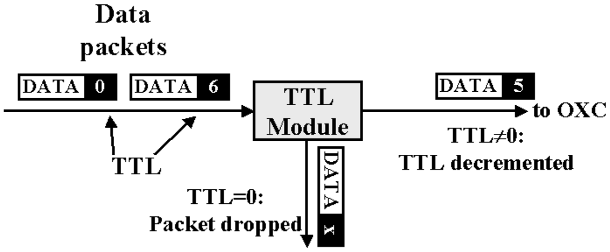

Fig. 1. A conceptual diagram of our optical time-to-live (TTL) decrementing module. A packet enters a switching node and first passes through the TTL module. If the TTL is nonzero, it is decremented by one and passes to the switch fabric. If the TTL is zero, it is dropped from the network and the resulting TTL value is irrelevant.

TTL field of the incoming packet is zero, the module processes the packet and generates a control signal that is used to drop the packet from the network. Our technique is independent of the length of the TTL and the packet, requires minimal control electronics, does not require the use of ultrashort RZ optical pulses, requires no guard time between the end of the TTL field and the rest of the packet, and has only a 2.4-dB power penalty at 10^{-9} bit-error rate (BER). In addition, while our module is experimentally demonstrated for 10 Gb/s NRZ-modulated packets, it can, with minimal modification, be made compatible with RZ-modulated data.

II. SYSTEM CONCEPT

A conceptual diagram of how our TTL module might operate within an optical switching node is shown in Fig. 1. A packet with a binary TTL field of known length (commonly 8 bs) enters the switching node and the TTL module. The TTL module checks the TTL of the incoming packet – if it is nonzero, the TTL is decremented by one and the packet passes through to the optical switch fabric. If the TTL of the incoming packet is zero, the packet is destroyed or dropped. In the zero-TTL case, the resulting TTL value after passing through the module is irrelevant as the packet is no longer within the network.

Optically decrementing the TTL by "1" requires optical implementation of a binary subtraction algorithm [3], [4]. One such method is decrementing-via-inversion. For any arbitrary starting binary value, subtracting "1" results in the least-significant-bit (LSB) being inverted (in binary subtraction, 1 minus 1 is 0, while 0 minus 1 results in a 1, after "borrowing" from the next higher column). If the LSB is a "1" bit, the subtraction is complete, as no borrowing is necessary. However, if the LSB is a "0" bit, we must reduce the next higher column by "1" as a result of the borrowing process (and if that next column is "0", leads to additional borrowing). This process is illustrated in Fig. 2(a). To decrement by one, invert each bit, beginning with the LSB, until a "1" bit is encountered. Invert the "1" bit, and then stop (as inverting the "1" bit completes the subtraction). All bits beyond the first "1" bit remain unchanged. This process results in a set of bits being replaced by their conjugates, and thus this binary subtraction can take place by replacing bits, up-to-and-including the first "1" bit, with their conjugates, as shown in Fig. 2(b) – the TTL data to be decremented in this figure consists of the bit sequence "11 010 000" ($\mathrm{MSB} \rightarrow \mathrm{LSB}$). To create the decremented output, the first 5 bits (starting with the LSB, up to and including the first "1" bit) are replaced with their conjugate (the shaded part of "databar"), while the rest of the bits remain the same (the shaded part of "data"). To implement this "decrementing-via-inversion" process in the optical domain, there are three main requirements: 1) a method of generating a set of conjugate data, 2) a fast method to replace data with its conjugate, and 3) a way to detect the first "1" bit in the TTL so the replacement process can be terminated.

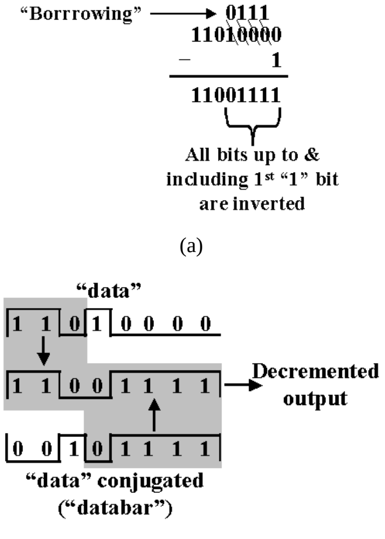

(a)

(b)

Fig. 2. (a) Binary subtraction of "1" from the 8-b binary value "11 010 000". Due to the borrowing process inherent to subtraction, the result, "11 001 111" has all bits, up to and including the first "1" bit, inverted. (b) Subtracting one from a binary number can be performed by replacing a series of bits (up to and including the first "1" bit) with their conjugates.

To generate a set of conjugate data to perform this process, we use gain saturation, or cross-gain modulation (XGM), in a semiconductor-optical amplifier (SOA) [5], [6]. An SOA is a well-studied optical device that has a number of interesting nonlinear properties, including XGM, cross-phase modulation (XPM), and four-wave mixing (FWM). A high-power data channel (the data to be conjugated) enters the SOA through one of its input ports, and a low-power continuous-wave (CW) signal enters via the other. An illustration of this process is shown in Fig. 3. In our experiment, both channels were at the same wavelength. When the data channel is "on" (a data "1" bit), the SOA provides the majority of its gain to the high-power signal, and the CW signal is squelched, resulting in a "0" bit on the CW wave at the output of the optical circulator. However, when the data channel is "off" (a data "0" bit), the CW signal receives most of the gain and a "1" bit results at the output. The resulting circulator output carries data conjugate to that of the data channel. The extinction ratio in this process can be $> 10\ \mathrm{dB}$. This inverted data can then be inserted where necessary to perform the "decrementing-via-inversion" procedure.

To replace data with its conjugate ("databar"), we use difference-frequency-generation (DFG) in periodically-poled lithium-niobate (PPLN) waveguide wavelength shifters [7]. These nonlinear devices use a cascaded $\chi^{(2)} : \chi^{(2)}$ process to efficiently shift an input signal to a new wavelength when a pump signal is present. The PPLN waveguide adds negligible spontaneous emission noise, operates with little to no chirp,

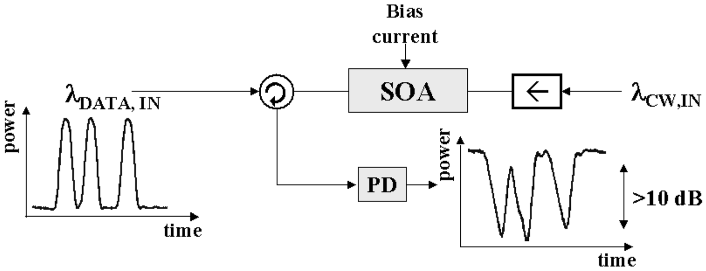

Fig. 3. Using gain saturation in a semiconductor-optical-amplifier (SOA) conjugate data can be generated at the output of the optical circulator. A high-power data channel and low-power continuous-wave (CW) channel (in our experiment, both channels were at the same wavelength) counter-propagate through the SOA. When the high-power channel is "on" (a "1" bit) the gain of the SOA is saturated and the low-power channel is strangled, resulting in a "0" bit. When the high-power channel is "off" (a "0" bit) the SOA contributes all its gain to the low-power channel, resulting in a "1" bit. The extinction ratio can be > 10 dB. Oscilloscope traces of the data sequence "10 100 100" are shown.

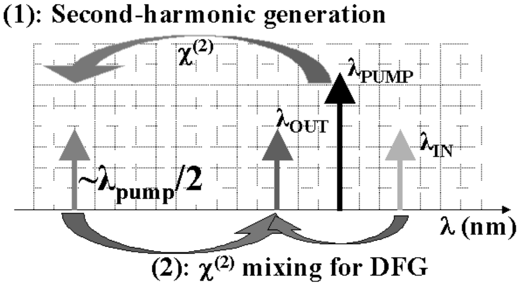

Fig. 4. Difference-frequency-generation (DFG) in a periodically-poled lithium-niobate (PPLN) waveguide results in "reflection" of an input wavelength $\lambda_{\rm IN}$ around the PPLN pump wavelength $\lambda_{\rm PUMP}$ via a cascaded $\chi^{(2)} : \chi^{(2)}$ process. First, via second-harmonic-generation, a channel at $\lambda_{\rm PUMP}/2$ is generated (1), and that channel mixes with the input $\lambda_{\rm IN}$ to produce $\lambda_{\rm OUT}$ (2).

has similar up- and down- conversion efficiency, induces negligible crosstalk at the output, and has a $>$ THz bandwidth [7]. In essence, the output wavelength is the input wavelength "reflected" around the pump wavelength, as shown in Fig. 4. These devices have seen application in a number of optical switching architectures [8]. A PPLN waveguide only λ-shifts an input channel to a new wavelength when a pump channel is present. By using two PPLN waveguides (one for the "data" and another for the inverted data, or "databar") and modulating the PPLN pumps appropriately, we choose when the "data" and "databar" signals are λ-shifted. Through the use of conjugate pump signals (so only one of the two signals is being λ-shifted at any given time), we can shift either "data" or "databar" as needed. By only λ-shifting "databar" to a new wavelength when the TTL-decrementing algorithm requires inverted data, and shifting "data" to a new wavelength at all other times (during the rest of the TTL, and during the rest of the packet), we can combine the resulting two λ-shifted signals and filter out the new wavelength, resulting in a combined, λ-shifted, and TTL-decremented channel.

III. EXPERIMENTAL SETUP

The experimental setup of our optical TTL module is shown in Fig. 5. An incoming data packet with an 8-b TTL field (with a 1-b guard time before the TTL to ensure proper timing, but no guard time between the end of the TTL and the rest of the packet data) is amplified and split into three branches. The middle branch is the "data" sent to the "data" PPLN waveguide. The top branch is transmitted through a 3-port optical circulator and through the SOA to create "databar" via gain saturation, as described in the preceding section. The final branch of the input data stream is sent to a receiver to check for the first "1" bit within the TTL. This is done by using this signal (after an OR gate) as the clock input to a D-flip-flop (DFF) that selects which of the data streams ("data" or "databar") is λ-shifted at any given moment by the PPLN waveguides by modulating the PPLN waveguide pumps. The DFF is the key to controlling the TTL-decrementing algorithm, with the "Q" output of the DFF controlling the "databar" PPLN pump modulator and the "$\rm Qbar$" output controlling the "data" PPLN pump modulator. This guarantees that only one of the two data streams will be λ-shifted at any given moment. A "TTL start" pulse (which we generate artificially, but could be provided by some sort of optical preamble detection such as in [9]) signals to the electronics the start of the TTL field and acts as both an input and a clock (after the OR gate) to the DFF, setting the "Q" output of the DFF to "1." As the "Q" output of the DFF is tied to the "databar" PPLN pump, the "databar" channel is shifted to a new wavelength, $\lambda_{\rm OUT}$, while the "Q" output remains high. The "data" channel, whose PPLN pump is controlled by the "$\rm Qbar$" output of the DFF, is not λ-shifted while the "Q" output remains high. We shift "databar" to $\lambda_{\rm OUT}$ only when the TTL-modification requires the conjugate data at the output – that is, until the first "1" bit is present in the TTL. The first "1" bit in the TTL of the packet acts as a "clock" input to the DFF, resetting the "Q" output to zero, shutting off the pump to the "databar" PPLN and turning on the pump to the "data" PPLN. This begins λ-shifting of the "data" stream to $\lambda_{\rm OUT}$. After filtering out all wavelengths save $\lambda_{\rm OUT}$, the two synchronized λ-shifted signals ("data" and "databar," which are never λ-shifted to $\lambda_{\rm OUT}$ simultaneously) are combined via a coupler and the result is a TTL-modified $\lambda_{\rm OUT}$ output channel.

The "Q" signal is also used to control a switch that drops the packet when the TTL of the incoming packet is zero. Any "1" bit within the TTL field will cause the "Q" output to reset to "0." Thus, if the "Q" output is still "1" after the TTL has passed, then there cannot be any "1" bits within the TTL field – it must have a zero value. A simple decision circuit is used to test "Q" immediately after the TTL has passed through the module. Only if all TTL bits are zero will the "Q" output remain high,

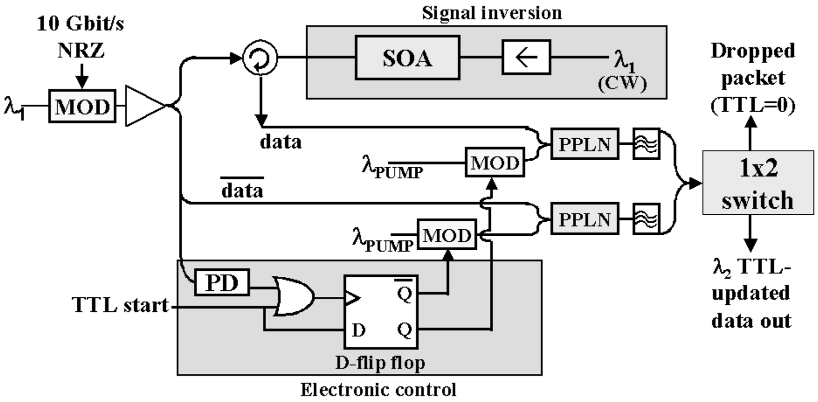

Fig. 5. The experimental setup of our optical TTL module. The initial data stream is split into three branches – one is used to create conjugate data ("databar") via gain saturation in an SOA, one is used as the data stream, and the last is detected and used to control the module. A D-flip-flop controls the insertion of conjugate data by controlling the PPLN waveguide pumps. The "TTL start" pulse is perhaps generated by an existing preamble-detection module, such as described in [9].

and thus a positive decision at this stage can be used to reset the DFF and drive an optical switch that drops the packet. While it is possible that some corruption of packet data can take place if the TTL is zero (as the "Q" output remains high, and thus "databar" is shifted to λ_{OUT} after the TTL has passed – it is for this reason that the decision circuit also resets the DFF), as packets are dropped or destroyed when the TTL of an incoming packet is zero, the resulting packet data may be irrelevant. In addition, as any packet with a nonzero TTL will stop shifting "databar" and begin shifting the "data" stream by the end of the TTL field, no guard time between the end of the TTL and the rest of the packet is required.

IV. RESULTS AND DISCUSSION

Three 20-B packets, each with varying TTL values and payload data, 0.8 ns of guard time between packets, and no guard time between the end of the TTL and the rest of the packet, are modulated at 10 Gb/s onto a 1545 nm signal wavelength (λ_{IN}) and sent into the TTL module. Two packets had nonzero TTL values, and one packet had a zero TTL. A closer look at the TTL field of one of the packets is shown in Fig. 6 for a packet with a TTL value of "00 111 100" (from the LSB to the most-significant-bit, or MSB). The input "data" signal for this packet is shown in Fig. 6(a), and the "databar" signal generated via gain saturation in the SOA (with a 1545 nm CW input to the SOA) is shown in Fig. 6(b). The SOA bias current was 80 mA. The first three bits of the TTL (starting with the LSB) are "001", so the first "1" bit is the third bit of the TTL. Thus, the "decrementing-via-inversion" algorithm should require that the first 3 bits of the TTL be replaced by their conjugates. The PPLN pump modulation signals (the outputs of the DFF) are shown in Fig. 6(c) and (d), where (c) is the "Qbar" signal modulating the "data" PPLN pump, and (d) is the "Q" signal modulating the "databar" PPLN pump. The pump wavelength for both PPLN's was 1550.12 nm, and the pump power into each PPLN was approximately 10 dBm, resulting in a λ-shifted output signal at 1555.24 nm at around $-13\ \text{dBm}$. The PPLN spectrum showing the input, pump, and λ-shifted output is shown in Fig. 7. After amplification and filtering to remove all wavelengths save 1555.24 nm, the resulting streams were combined and the TTL-modified result is shown in Fig. 6(e). While our technique results in a shift of the input wavelength to a new output wavelength, this can be alleviated by using a conventional wavelength shifter (or a second PPLN with a

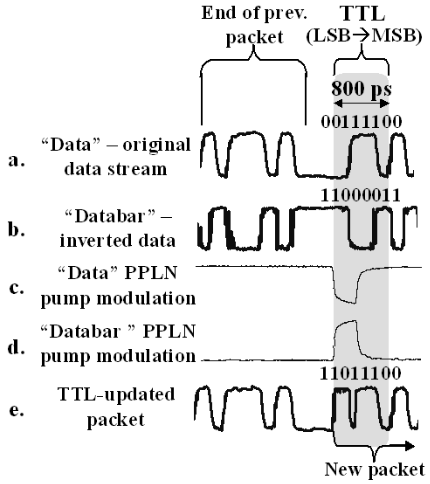

Fig. 6. A close-up look at the TTL field of a packet with a TTL value of "00 111 100" ($\text{LSB} \rightarrow \text{MSB}$) (a) The TTL field of the "data" stream. (b) The TTL field of "databar", the conjugate data generated via gain saturation in the SOA, showing a TTL value of "11 000 011". (c) The control signal for the "data" PPLN pump – "data" is λ-shifted at all times save from the start of the TTL until after the first "1" bit in the TTL. (d) The control signal for the "databar" PPLN pump – "databar" begins λ-shifting when the TTL starts and ends immediately after the first "1" bit. (e) A close-up view of the TTL-modified output, with a new TTL of "11 011 100" (the bits up to and including the first "1" bit having been replaced with "databar").

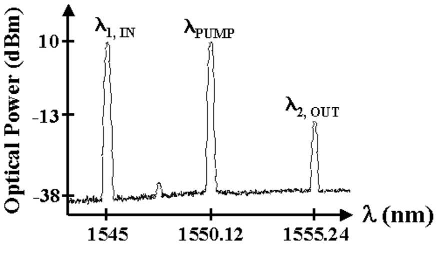

Fig. 7. The optical spectrum at the output of the "data" PPLN waveguide, showing the three signals – "data" at $\lambda_{\text{IN}} = 1545\ \text{nm}$, the λ-shifted output at $\lambda_{\text{OUT}} = 1555.24\ \text{nm}$, and the PPLN pump at 1550.12 nm.

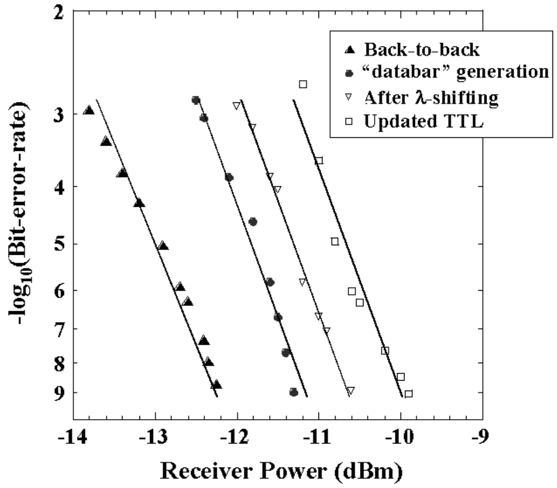

Fig. 8. Power penalty curves for the optical TTL module. The total penalty is ~ 2.4 dB.

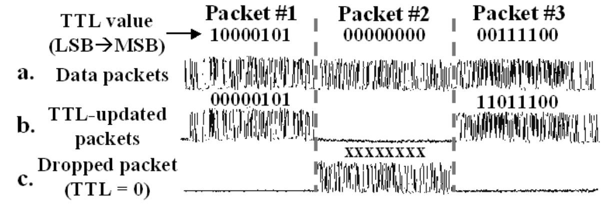

Fig. 9. (a) A number of input packets to the TTL module, each with different TTL values and data payloads. (b) Packets with nonzero TTL values have their TTL decremented and then pass through the module. (c) When a packet with a TTL value of zero enters the module, a signal is generated that drops the packet from the network (in this case using an optical switch), and the resulting TTL value is irrelevant.

pump wavelength identical to the first, causing a second shift around the pump back to the original input wavelength).

Power penalty curves are shown in Fig. 8. The total power penalty of our module is ~ 2.4 dB when compared to the back-to-back receiver sensitivity at 10^{-9} bit-error-rate, arising mostly by the λ-shifting in the PPLN and insertion of parts of "databar" into the TTL field. Using the current setup, it may be difficult to traverse a large number of hops, and as such this system may be viable only in a low-hop network core. However, the use of a higher-efficiency PPLN waveguide and an SOA with lower chirp (or an optical 2/3 R regenerator) may reduce these sensitivity penalties considerably, allowing for transmission beyond a handful of nodes.

To demonstrate packet dropping, a lithium-niobate optical switch was placed in series with the TTL module and controlled by a decision circuit as described in the above section. The circuit checks the "Q" value on the first bit after the TTL ends – if it is high, it signals the switch to drop the packet and resets the DFF so that the "Q" signal is again "0." The three packets are shown in Fig. 9(a), and the switch outputs are shown as Fig. 9(b) and (c), where (b) shows the "through" port and (c) the "drop" port.

While this technique was demonstrated for NRZ-modulated data, this technique can be applied to RZ systems without major modification. By replacing the low-power CW signal input to the SOA with a constant stream of RZ "1" bits synchronized to the input, a "databar" stream of RZ bits can be generated. The rest of the system requires no modification—such an implementation also removes the need for the 1-bit guard time prior to the start of the TTL. Without further investigation as to the effects of the PPLN and SOA on other types of signals (CRZ, CSRZ), we are unable to say whether this technique can be extended to other modulation formats.

V. SUMMARY

"Routing loops," where a misdirected packet is routed in circles indefinitely, can be a significant problem in modern-day networks, leading to severe network congestion if misrouted packets are not quickly dropped from the network. To combat this problem, IP packets incorporate a "time-to-live" field that is decremented by one at each "hop" within the network – when the TTL value reaches zero, the packet is dropped from the network. To begin a move toward all-optical routing within the network core, it may be necessary to implement this TTL decrementing process in the optical domain.

We demonstrate an optical TTL decrementing module that acts upon a standard 8-b binary TTL field located within an optical packet. If the value of the TTL field within the packet is nonzero, the TTL is decremented via a decrementing-via-inversion process and the packet passes through the module, perhaps toward an optical switch fabric. If the TTL field is zero, however, the module generates a signal that can be used to drop or destroy the packet. Our technique uses gain saturation within an SOA to generate conjugate data, and difference-frequency-generation in a PPLN to enable wavelength conversion and insertion of conjugate data. This technique is independent of the length of the TTL and packet, requires minimal electronic control, requires no guard time between the end of the TTL field and the rest of the packet, and has only a 2.4-dB power penalty at 10^{-9} BER.

REFERENCES

[1] V. Paxson, "End-to-end routing behavior in the internet," *IEEE/ACM Trans. Networking*, vol. 5, pp. 601–615, Oct. 1997.

[2] W. Hung, K. Chan, L. K. Chen, C. K. Chan, and F. Tong, "A routing loop control scheme in optical layer for optical packet networks," in *Conf. Opt. Fiber Commun. (OFC) 2002*, 2002, ThGG111, pp. 770–771.

[3] R. E. Fowkes, "Hardware efficient algorithms for trigonometric functions," *IEEE Trans. Computers*, vol. 42, pp. 235–239, Feb. 1993.

[4] H.-I. Jeon and M. A. G. Abushagur, "Digital optical arithmetic processor based on symbolic substitution," in *Proc. 20th Southeastern Symp. Syst. Theory (1998)*, 1998, pp. 221–223.

[5] P. Parolari, L. Marazzi, M. Connen, and M. Martinelli, "SOA-based all-optical threshold," in *Conf. Lasers and Electro-Optics (CLEO) 2000*, 2000, CWK29, pp. 309–310.

[6] M. C. Cardakli, D. Gurkan, S. A. Havstad, and A. E. Willner, "Variable-bit-rate header recognition for reconfigurable networks using tunable fiber-Bragg-gratings as optical correlators," in *Conf. Optical Fiber Commun. (OFC) 2002*, vol. 1, 2000, TuN2, pp. 213–215.

[7] K. R. Parameswaran, J. R. Kurz, R. V. Roussev, and M. M. Fejer, "Observation of 99% pump depletion in single pass SHG in a PPLN waveguide," in *Conf. Lasers and Electro-Optics (CLEO) 2001*, 2001, CFB3, p. 549.

[8] G. S. Kanter, P. Kumar, K. R. Parameswaran, and M. M. Fejer, "Wavelength-selective pulsed all-optical switching based on cascaded second-order nonlinearity in a periodically poled lithium-niobate waveguide," *IEEE Photon. Technol. Lett.*, vol. 13, pp. 341–343, Apr. 2001.

[9] 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.

J. 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. He is currently working toward the Ph.D. degree at USC.

He is currently a Research Assistant in the USC Optical Communications Laboratory under the direction of Alan E. Willner. His research interests focus on all-optical signal processing and implementation of routing and networking functions in the optical domain, as well as signal conditioning and monitoring, Raman amplification, and optical code-division-multiple-access (OCDMA) systems. He is the author and/or coauthor of more than 20 papers in technical journals and conferences in the field.

Mr. McGeehan is a Member of the Lasers & Electro-Optics Society (LEOS) and Education societies.

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 working toward 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 at USC. His research focuses on all-optical signal processing to enable implementation of networking functions in the optical domain.

Mr. Kumar is a Member of the Eta Kappa Nu academic honor society and a Student Member of the Optical Society of America (OSA).

Deniz Gurkan (S'92) received the B.S. and M.S. degrees in electrical engineering from Bilkent University, Ankara, Turkey, in 1996 and 1998, respectively, and the Ph.D. degree in electrical engineering from the University of Southern California (USC), Los Angeles.

She is currently a Postdoctoral Researcher at the Information Sciences Institute, USC, Marina Del Rey, and a Lecturer at the California State University Long Beach. Her research is on implementation of Internet routing functions, such as routing, flow control, and switching in the optical domain in order to enable all-optical wavelength-division-multiplexed packet switching networks.

Dr. Gurkan is a Member of the IEEE Communications Society and the Optical Society of America (OSA).

S. M. R. Motaghian Nezam (S'03) received the B.S. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1995 and the M.S. degree from the University of Southern California (USC), Los Angeles, in 2002. He is currently working toward the Ph.D. degree at USC.

In 2000, he joined the Optical Communications Laboratory at USC. His main research interests include optical performance monitoring and equalization (chromatic and polarization-mode dispersion monitoring and equalization), data modulation formats, analog transmission systems, and long-haul wavelength-division-multiplexing (WDM) transmission. He is the author and/or coauthor of more than 30 papers in international journals and conference proceedings.

Mr. Nezam was a recipient of a Lasers & Electro-Optics Society (LEOS) Graduate Fellowship Award in 2003. He is a Student Member of LEOS, the IEEE Communications Society (ComSoc), and the Optical Society of America (OSA).

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.

K. R. Parameswaran (M'03) received the B.Eng. degree in electrical engineering from Carleton University, Ottawa, ON, Canada, in 1995 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1997 and 2002, respectively. His dissertation work involved the development of highly efficient nonlinear optical frequency mixers for use in communications, spectroscopy, and quantum optics.

He is currently working as a Postdoctoral Researcher in fiber-optic sensors at Stanford University. His research interests include photonic devices, microfabrication, and remote sensing.

Dr. Parameswaran is a Member of the Optical Society of America (OSA).

M. M. Fejer (M'03), photograph and biography not available at the time of publication.

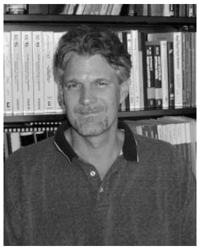

J. Bannister (S'80–M'80–SM'95) received the B.A. degree (with high distinction) in mathematics from the University of Virginia, Richmond, in 1977 and the M.S. degree in electrical engineering and the M.S. and Ph.D. degrees in computer science from the University of California, Los Angeles, in 1980, 1984,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

Joseph D. Touch (S'83–M'92–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 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 working on fault-tolerant networks (DynaBone, NetFS, X-Tend), optical Internet wide-area networks and local-area networks (LANs) (POW, OCDMA), geographic routing (GeoNet), and Smart Space devices for user presence. He has led projects ranging from gigabit LANs (ATOMIC2), NIC design (PC-ATOMIC), and multicast web caching (LSAM), 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 and now runs a program for summer students (SGREP) and advises a number of graduate students. He has several patents current and pending and is author of a number of publications, including the book *High Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication*, and he serves on the editorial board of Elsevier's *Computer Networks*. 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 and the Association of Computing Machinery. He Co-Chairs the IEEE ITC (Internet) committee and is active in the IETF. He also serves on the editorial board of IEEE NETWORK and is a Member of several program committees, including IEEE INFOCOM (since 1994) and Sigcomm, and was Vice General Chair of Opticomm 2001.