Physical Layer Scalability of WDM Optical Packet Interconnection Networks

Odile Liboiron-Ladouceur, *Student Member, IEEE*, Benjamin A. Small, *Student Member, IEEE*, and Keren Bergman, *Member, IEEE*, *Fellow, OSA*

***Abstract*—The physical layer scalability of a packet-switched optical interconnection network utilizing semiconductor optical amplifier (SOA) switch elements is investigated experimentally and with numerical modeling. Optical packets containing payloads of multiple wavelength-division-multiplexing (WDM) channels are propagated through cascaded SOA-based switching nodes in a recirculating test-bed environment. Experiments show that bit-error rates (BERs) below 10^{-9} can be maintained through 58 switching nodes for the entire eight-channel 10-Gb/s-per-channel payload distributed over 24.2 nm of the C-band. When the packet payload consists of a single 10-Gb/s channel, 98 node hops can be traversed before a BER of 10^{-9} is exceeded. In conjunction with the experiments, a novel phenomenological modeling technique is developed in order to forecast the scalability of SOA-based WDM packet interconnection networks. This technique is shown to yield results that correlate well with the experimental data. These investigations are presented as predictors of the physical limitations of large-scale WDM packet-switched networks.**

***Index Terms*—Interconnection networks (multiprocessor), optical interconnections, packet switching, semiconductor optical amplifiers.**

I. Introduction

OPTICAL interconnection networks are a promising means of routing high-bandwidth optical packets in applications ranging from data communications and storage to high-performance computing. As communication between numerous interconnected elements drives the increasing demand for cross-sectional bandwidth, optical interconnection networks can better address these challenges by providing throughput scalability [1]. Maintaining routed packet traffic in the optical domain enables full exploitation of the bandwidth provided by dense wavelength-division multiplexing (DWDM) [2], [3]. As data rates increase, the tremendous amount of bandwidth offered by DWDM technology can potentially reduce the cost per bits and the power consumption compared to copper-only communication solutions [4].

The overall performance and scalability of packet-switched optical networks have been studied for a range of architectures [5]–[11]. A critical measure of the feasibility of packet-switched networks intended for implementation in the optical domain is the physical layer scalability. Although the topology may scale in terms of throughput and latency, it does not necessarily follow that the optical physical layer also scales in terms of maintaining end-to-end signal integrity without requiring regeneration. In this paper, we examine the physical layer performance both experimentally and with numerical modeling, specifically for a WDM optical packet interconnection network based upon the data vortex architecture and implemented with semiconductor optical amplifier (SOA)-based switching elements. Since SOAs are the predominant switching elements employed in optical packet networks, our results may easily be extended to other similar architectures [12]–[14].

The potential benefits of SOA-based switching elements are well known and include fast switching times, high extinction ratios, sizable operating gain bandwidth, and a relatively compact footprint [2], [15], [16]. SOAs, however, also necessarily introduce amplified spontaneous emission (ASE) noise that mixes with the propagated optical signal as a side effect of the amplification process. The noise figure is manifested in the measured degradation of the optical signal-to-noise ratio (OSNR) as packets propagate through a cascade of switching nodes [17]. Additional effects must be considered when WDM signals are propagated. The non-constant gain profile of the SOA across the C-band will cause each WDM payload channel to experience a slightly different gain. The finite spectrum profiles of passive components used in the switching nodes are also generally not uniform. The combination of these effects is accentuated with increasing number of cascaded nodes as each payload channel spanned across the C-band experiences different gain and loss propagating through the entire interconnection network.

For large-scale packet-switched networks, the mean number of intermediate nodes or hops each packet traverses will generally scale logarithmically with the number of ports [5], [11]. Thus, the physical layer scalability of the data vortex as well as other SOA-based interconnection networks is ultimately determined by the accumulation of these signal degradation phenomena as the packet propagates through a cascade of many switching nodes.

In this work, the propagation of optical packets carrying multiple 10-Gb/s WDM channels payloads through a cascade of switching nodes is investigated experimentally in a recirculating loop test-bed environment. Previously, recirculating loop approaches have been utilized in the context of cascaded SOA booster amplifiers for WDM long-haul transmission systems [18], [19], specifically to investigate various phenomena such as polarization-dependent gain (PDG) [20], cross-gain modulation

Manuscript received November 30, 2004; revised September 14, 2005. This work was supported in part by the National Science Foundation under Grant ECS-0322813 and ECS-0532762 and by the U.S. Department of Defense under Subcontract B-12-644.

The authors are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: ol2007@columbia.edu; bas@ee.columbia.edu; bergman@ee.columbia.edu).

Digital Object Identifier 10.1109/JLT.2005.859852

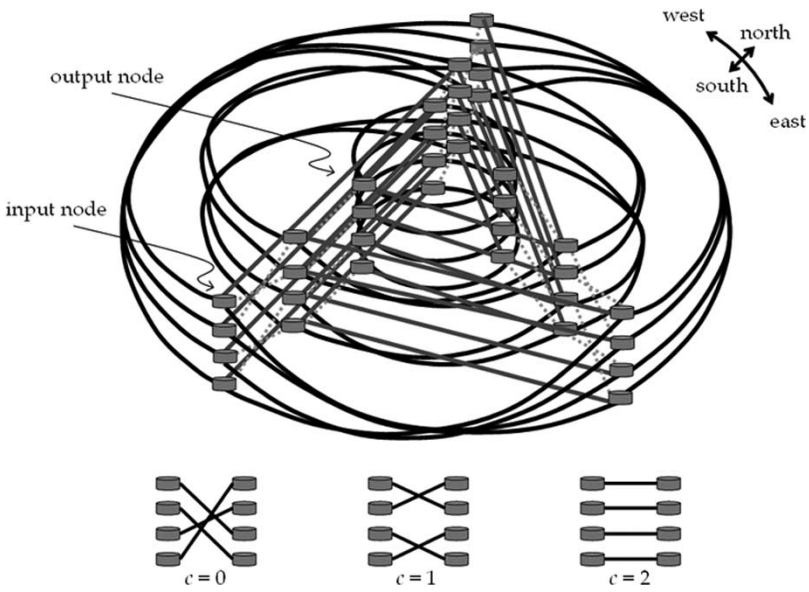

Fig. 1. Data vortex topology with A = 3, H = 4, C = 3 corresponding to a 12-port interconnection network that contains 36 switching nodes.

(XGM) [21], and bandwidth efficiency [22]. The SOA functionality within WDM optical packet switching (OPS) networks, however, significantly differs from its operation in transmission systems. For typical WDM transmission systems, SOAs are designed to operate primarily as high-saturation gain elements separated by kilometers of fiber. In the context of an SOA-based recirculating loop simulating transmission system, the maximum number of cascaded SOA gates was found to be 11 using holding light injection [23]. In OPS networks, the primary function of the SOA is to efficiently route a broadband packet and compensate for small switching node losses. For this application, SOAs are separated by only a couple of meters and are expected to operate in a linear low-gain regime, resulting in minimal added noise and signal distortion. It is expected that these operating conditions can enable the cascading of large numbers of switch elements and, thus, the scalability of OPS networks.

In this paper, we report on experiments and simulations that demonstrate the successful cascading of SOA-based OPS gates. The SOA switching nodes are constructed in accordance with the data vortex network architecture. Our experiments show that a bit-error rate (BER) of 10^{-9} can be maintained for 58 hops with an eight-channel 10-Gb/s-per-channel payload spanning 24.2 nm of the C-band. When the packet payload consists of a single 10-Gb/s channel, it is shown that 98 node hops can be traversed before exceeding a BER of 10^{-9}. In conjunction with the experiments, a novel phenomenological modeling technique is developed in order to capture the effects of cascaded SOAs on packet signal integrity.

II. Data Vortex Overview

A. Data Vortex Network Topology

OPS networks offer a promising means to efficiently provide ultra-high bandwidth communications between thousands of interconnected elements [1], [24]. There are, however, several key challenges to implementing viable large-scale interconnection networks in the optical domain. Perhaps the most significant is the lack of adequate random-access optical buffering, particularly for hybrid data structures such as WDM [25], [26]. Optical signal processing is also difficult to realize with low latencies [27]. Due to these fundamental shortcomings, it is imperative that traditional architectures used for electronic switching fabrics not be simply transplanted to the optical domain but that the characteristics of photonic technologies be considered in the development of OPS network architectures.

The data vortex [28] is a self-routing deflection network topologically scalable for large (e.g., 10k-port) interconnections and specifically designed for implementation with photonic technologies [29]. Its deflection routing mechanism eliminates the need for internal optical buffering while meeting large-scale interconnection networks' key requirements of low switching latency and high throughput. In this time-slotted network architecture, the payload data are encoded on multiple DWDM channels in parallel with the header and frame signal wavelengths. One packet at a time is processed within each node in a particular timeslot. The implementation of the header and the frame is at a lower bit rate than the payload to facilitate the electronic decoding and interpretation of the destination information. A particular data vortex network can be described by the three parameters C, A, and H, corresponding to the cylinder, angle, and height parameters, respectively. The system illustrated in Fig. 1 corresponds to a switch fabric supporting 12 input and 12 output ports [30].

In the schematic representation of Fig. 1, packets ingress from input nodes at the outermost cylinder and exit from output nodes at the innermost cylinder. Upon entering a node, the packets are routed based on the destination address encoded in the header signal and the deflection signaling from an adjacent node. If the address does not match the user-programmed value, or if the next node is busy, the packet is deflected. If

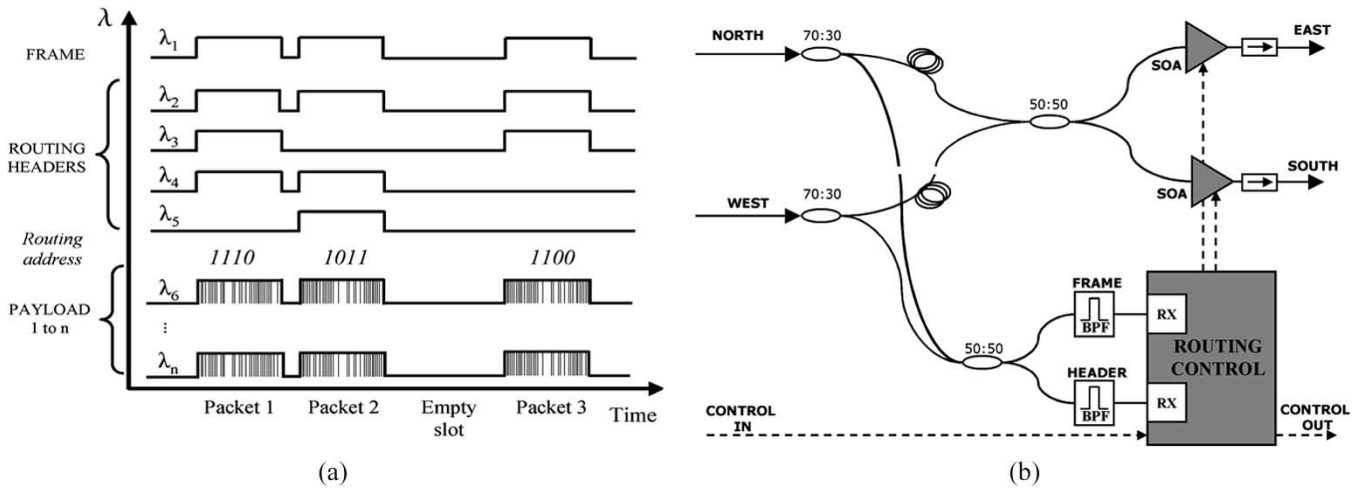

Fig. 2. (a) Packet structure: payload and routing headers encoded on DWDM channels. (b) Switching node configuration of the data vortex network (BPF: optical bandpass filter; RX: optical receiver module; plain line: optical signal; dashed line: electrical signal).

the destination output buffers are busy, the innermost cylinder allows the packet to circulate.

B. Switching Node Architecture

The switching nodes within the data vortex architecture are implemented with conventional fiber-optic technologies. As packets propagate through a node, the address and frame routing signals, which are transmitted within the packet [Fig. 2(a)], are decoded at each switching node. The simple switching node structure for this self-routing deflection network is shown in Fig. 2(b). Each node has two input ports labeled west and north, and two output ports labeled east and south. When the header address matches the user-programmed value, the packets are directed to the south output port. When the address does not match or when the next node is occupied, the packets are deflected to the east output port. The electronic control circuitry operates at the packet rate, which is on the order of megabits per second. At each input port, a small portion of optical power is tapped off using a 70/30 coupler and the header and frame bit information are isolated using bandpass filters. Their respective bits are converted to electronic signals using a commercial optical receiver module. The address information is processed with the electronic input deflection signal. The payload data are transparent to the internal node and delayed while the routing decision is processed. SOAs are used as switching elements to appropriately route the packet either to the downstream node (south) or to the adjacent node (east). A deflection signal is also generated as necessary to ensure that packets avoid collision. The data vortex architecture is explained in more detail in [29]–[32].

Simulations of the data vortex as a large-scale switching fabric have shown that for a heavily loaded 10k $\times$ 10k port data vortex implementation, 99.999% of the injected packets propagate through fewer than 58 internal switching nodes. The non-Gaussian distribution curve exhibits a median hop count of 19 (Fig. 3). For switching networks in general, a greater

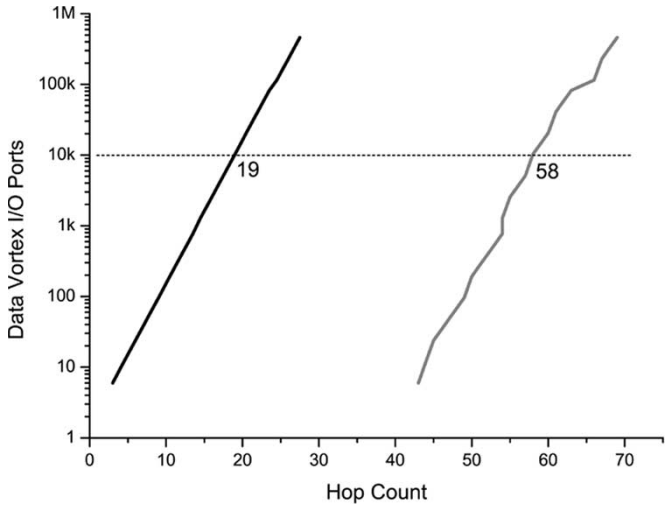

Fig. 3. Data vortex I/O port size versus hop count under heavy load. Network size for the median number of hops (black) and for the 99.999th percentile of the injected packets (gray).

number of hops correspond to a larger input and output port count, proportional to $\log(N)$ for an $N \times N$ port network.

III. Recirculating Loop Setup

A. Experimental Recirculating Test-Bed

A schematic of the test-bed constructed to experimentally emulate the propagation of WDM optical packets through multiple data vortex switching nodes is shown in Fig. 4. Eight temperature-controlled DFB lasers with polarization controllers are combined by an 8:1 planar coupler into a booster SOA, which compensates for transmission losses. The WDM channels are distributed across the C-band ranging from 1530 to 1575 nm. The channels are simultaneously modulated with a 10-Gb/s nonreturn-to-zero (NRZ) pseudorandom bit sequence (PRBS) of length $2^9 - 1$ produced by the fast pattern generator (PPG1) via a single $LiNbO_3$ modulator. All eight channels are then decorrelated by 450 ps/nm using 25 km of single-mode

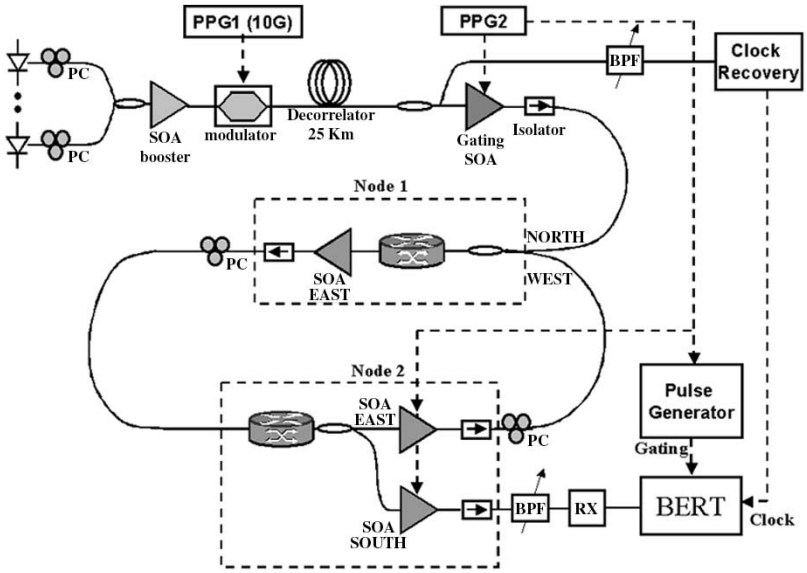

Fig. 4. Experimental setup of the switching-loop test-bed (PPG1: fast programmable pattern generator; PPG2: slow programmable pattern generator; PC: polarization controller; SOA: semiconductor optical amplifier; boxed arrow: optical isolator; BPF: bandpass filter; MOD: $LiNbO_3$ modulator; RX: optical receiver module; BERT: BER tester).

fiber. A 25.6-ns-long packet containing eight WDM payload channels, each containing 256 bits, is created using a separate SOA to gate the continuous data stream.

Two modified data vortex switching nodes are placed in a 64-ns-long recirculating loop. The header and frame information are not transmitted along with the packets, but instead provided directly to the switching nodes from the slow programmable pattern generator (PPG2), bypassing the optical-to-electrical conversion and the routing control logic gates. The modification is suitable in the context of the recirculating loop since the measurement addresses the number of node hops traversed by the packets, assuming proper routing logic function.

A single packet is injected into the loop though the north input port of node 1. A packet propagating in one loop undergoes two node hops. After a predetermined number of hops, the packet drops out of the loop through the south output port of node 2 at a selected time interval. The east output port of node 1 is always selected by the control signaling to ensure that packets remain in the loop. The packet can then be kept in the loop by selecting the east port of node 2 or ejected through the south port of node 2.

The SOAs are commercial devices (Kamelian OPB-10-10-X-C-FA) with a noise figure of 6.5 dB, an unsaturated gain of 8 dB, and a saturation input power of approximately 0 dBm for an operating current of 40 mA. The SOA gain peak is centered at 1465 nm by design to achieve lower ASE noise within the C-band. In each node, the SOA gain compensates the losses accrued by the couplers and connectors to achieve nearly no net loss through the entire loop. Polarization controllers are placed after each node to compensate for the 0.2-dB PDG of the SOAs. The optical receiver follows a narrow bandpass filter that selects the WDM channel to be measured. The last SOA seen by the packet is at the south output port of node 2, which is set to provide enough gain to reach the sensitivity limit of the optical receiver. The electronic payload information is then directed to the sampling scope and BER tester (BERT).

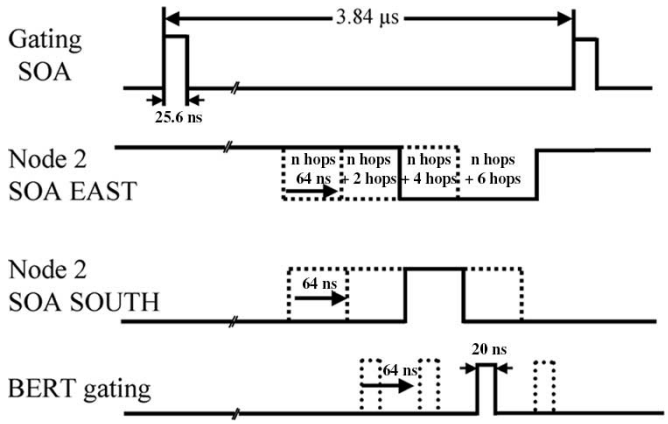

Fig. 5. Packet synchronization. The SOA control signals and the gating signals are delayed by 64 ns to increment the propagation by one loop. One loop consists of two hops. SOA SOUTH is switched ON to drop the packet for BER measurement while SOA EAST is switching OFF.

B. Packet Synchronization Methodology

As shown in the timing diagram in Fig. 5, a single 25.6-ns packet is injected once every 3.84 μs (60 loop roundtrips) into the 64-ns-long recirculating loop. Thus, only one packet at a time propagates within the recirculating loop. The extra time between packets is used to clear out any accumulated noise prior to launching the next packet. When a packet exits the loop via the south port of node 2, the SOA in the east port is turned OFF to minimize any additional leaked noise during BER measurements. To increment the number of hops by two, i.e., one loop, the control signals to node 2 and the gating signal to the error detector are delayed by 64 ns (corresponding to one loop length).

BER measurements are performed on all eight WDM channels of the packet using the 10-Gb/s BERT, which is externally gated by a pulse generator synchronized with the programmable pattern generator that also controls node 2 and the gating SOA. The length of the gating signal generated by the pulse generator for the error detector is 20 ns. The gating pulse is centered in the middle of the 25.6-ns-long packet to frame the measurement of valid data outside the rise and fall times of the cascaded SOA switching elements. The time delay between the gating signal and the node control signal in Fig. 6 corresponds to the total latency between the south SOA of node 2 and the error detector.

As shown in Fig. 4, the clock reference used by the error detector is extracted directly from the continuous data stream of the WDM channel under measurement using the OC-192 clock recovery option of an optical sampling module with a 28-GHz bandwidth. This avoids the phase drift that may occur between the received data and the clock signal.

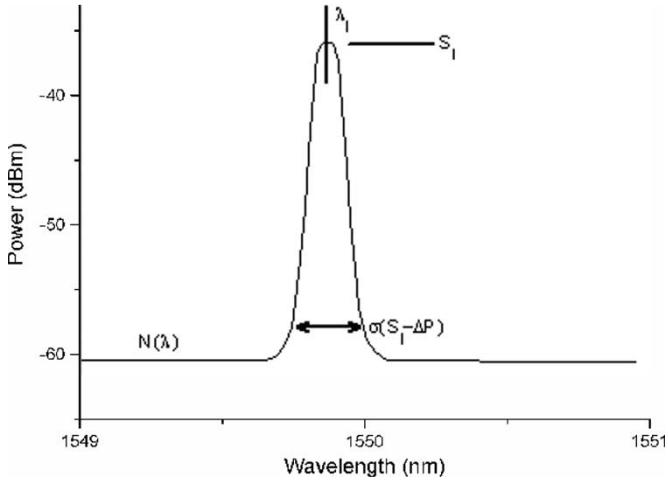

Fig. 6. Annotation of the spectrum parameters emphasized in the phenomenological model.

IV. Modeling

A. Data Vortex Network Topology

Traditional approaches for modeling SOAs are based on physical rate equation [16], [17], [33] or traveling wave analyses [16], [34]. These techniques have shown to be quite useful for simulating transmission and long-haul communication systems that utilize SOAs as amplifiers, accurately predicting amplification, noise figure, and even some nonlinear effects for single-channel systems with continuous data streams (when the SOA is operated as a high-gain amplifier) [18]. However, many of the effects that become important for low-power DWDM bursts are left as second-order consequences of the models and are hence very difficult to match accurately with empirical data in OPS networks. For example, conventional modeling techniques require a vast quantity of precise physical characteristics of the SOA's wavelength-dependent absorptivity and gain efficiency in order to make accurate predictions about the signal gain and gain saturation [16] for a particular operating wavelength. Furthermore, fabrication parameters like facet reflectivity might also be necessary for determining second-order device behavior [33]. Clearly, when these values are not available or when they are difficult to isolate and measure, the precision of the model suffers, especially when higher order effects are of interest. Thus, a numerical phenomenological model has been designed that concentrates on the properties important to low-power DWDM packet switching, and although it is based upon empirical observations instead of physical first principles, this model has shown predictive powers that surpass those of the conventional techniques when applied to the particular application of SOAs in OPS networks [34].

B. Methodology

This experiment investigates SOAs within the context of large switching networks and systems, and the characteristics of interest are system-level ones. Accordingly, the parameters that have been isolated as the most important are directly related to system-level functionality: noise figure, gain and gain compression, spectral broadening, and cross-gain and cross-phase modulation [17]. Many of these parameters are directly related to features of the signal power spectrum, while others require a more traditional time-domain perspective.

This empirically based modeling technique characterizes spectral and signal degradations due to the SOAs and couplers. It then makes projections about the evolution of these features over a large number of node hops. The specific features of interest are the peak signal power S_i and the peak width at a particular power level $\sigma_i(S_i - \Delta P)$ for each signal centered at wavelength λ_i, and also the shape of the noise floor $N(\lambda)$ (see Fig. 6). We quantify the modification of each parameter by the SOA and extrapolate these trends for large system diameters. In addition to these spectral features, time-domain characterizations of the signal waveform allow the model to incorporate carrier dynamics and other effects not manifested by spectral deformities. The time-domain parameters used are based upon carrier lifetimes, carrier densities, and quantum efficiency of the SOA, consistent with many of the conventional modeling techniques.

This methodology has been shown to make large system projections quite accurately when given the appropriate empirical information [35], [36]. Recent improvements in our techniques have further enhanced the modeling capabilities of WDM systems, specifically of the type currently under discussion.

C. Parameters

All of the spectra modeling parameters are based on empirical observations and data collected in experiments similar to the one discussed above. The necessary information like wavelength dependence of signal peak gain and spectral broadening can be extracted from the power spectra obtained in experiments that contain many fewer SOA hops. These data can be extrapolated to a system with similar devices operating under similar conditions, but with a much larger number of SOA hops. Thus, our model is predictive in that its phenomenological parameters are derived from smaller systems; it synthesizes this

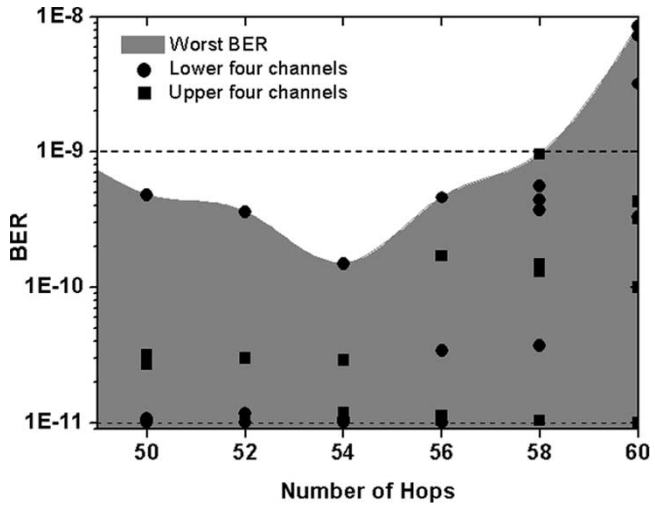

Fig. 7. BER for eight payload channels versus number of hops; dashed line for BER of 10^{-9}. The BER measurement of each payload channel falls within the gray region. Annotation of the spectrum parameters emphasized in the phenomenological model.

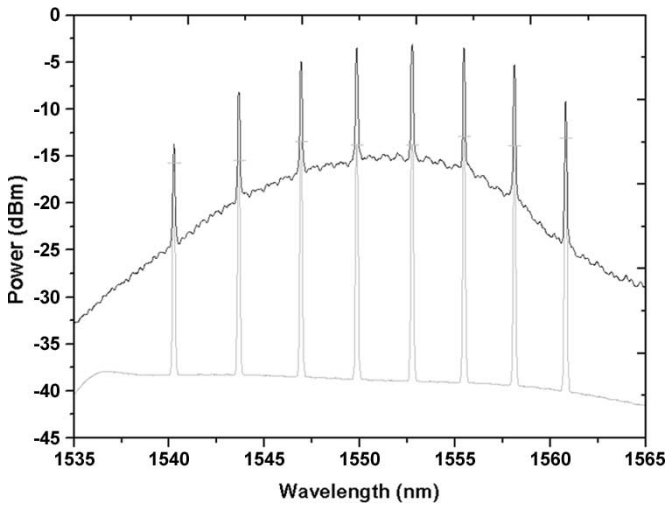

Fig. 8. Multichannel spectrum profile: the input packet spectrum (light gray) and after 58 node hops (dark gray) defined as the average power, taking into account the 0.67% duty cycle where one 25.6-ns-long packet is sent through the recirculating loop every 3.84 μs.

information and projects the behavior of larger systems, which may be difficult to realize experimentally.

The evolution of the noise floor required a thorough study of the SOA's ASE noise profile and of the loss profile of the couplers and passive optical components used; a simple parabolic fit was not sufficient. For smaller systems, these small nonuniformities would be almost unnoticeable, but for our larger system they become substantial. Almost certainly, a conventional model would have difficulty in capturing all of these important features, but this analysis ensures that no information that is important to system-level behavior is relegated to the second order.

For 10-Gb/s payloads, the spectral width is sampled at 15.6 dB below the signal peak (i.e., $\Delta P = 15.6$ dB). Empirical evidence suggests that for higher data rates, the value of ΔP used should be lower. Focusing the scope of the model to a specific modulation bandwidth limits its universality and flexibility, but allows for the current application to be studied more thoroughly. This tradeoff is consistent with the original goals in developing a predictive phenomenological model and yields some interesting and unique results about multichannel WDM packets when considered alongside empirical observations.

V. Results

A. Saturation–Sensitivity Tradeoff

Packets containing eight-channel 10-Gb/s-per-channel payload are measured to successfully propagate through 58 hops and maintain a BER below 10^{-9} on each payload channel (Fig. 7). The packet's eight-channel payload spectra profiles after two hops and after 58 hops are shown in Fig. 8. The input average peak power of each channel, distributed evenly across the C-band from 1540 to 1561 nm, varies between -12.9 and -15.7 dBm, and the total launched average input power per packet is -4.9 dBm, below the 0-dBm input saturation power of the SOAs. After 58 hops, the total optical output power reaches 2.7 dBm. As the packet progresses through multiple hops, the ASE noise rises and takes the form of a bell-shaped spectrum profile with a peak noise floor increase of 23.6 dB near the channel at 1552.8 nm. The gain declines for higher and lower channel wavelengths. Fig. 9 shows the corresponding measured and simulated eye diagrams. The eye diagram was predicted by the phenomenological modeling technique, as was the power spectrum with OSNR, which matched the empirical result to within approximately 0.2 dB.

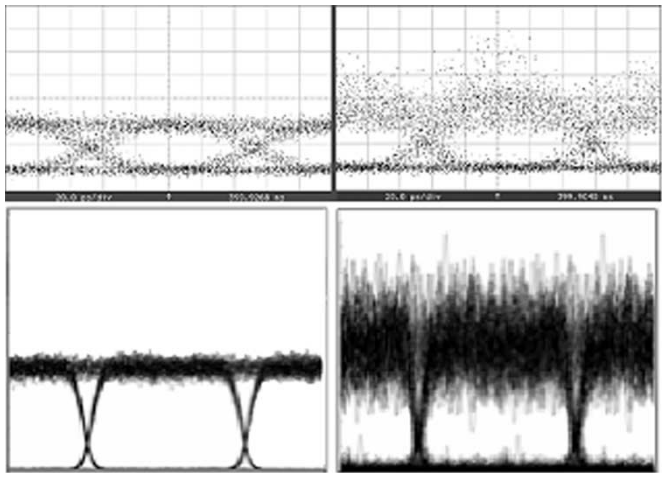

Fig. 9. Optical eye diagram of channel at 1558.10 nm after 2 and 58 hops. Simulated eye diagram of one channel after two hops (left) and after 58 hops (right).

The peak near 1553 nm (Fig. 8) is attributed to the cascading wavelength-dependent loss of the passive elements (couplers and isolators) in the loop test-bed. The ASE noise accumulation as the packet progresses through the recirculating loop leads to a total optical power increase causing the SOA to eventually operate in the saturation regime [16], [37]. Because the relatively fast response and gain recovery times of the SOAs are on the order of the payload data modulation speeds (i.e., $\sim$ 100 ps), XGM occurs as the SOAs approach operation in the saturation regime, and the BER of the payload data is degraded [38].

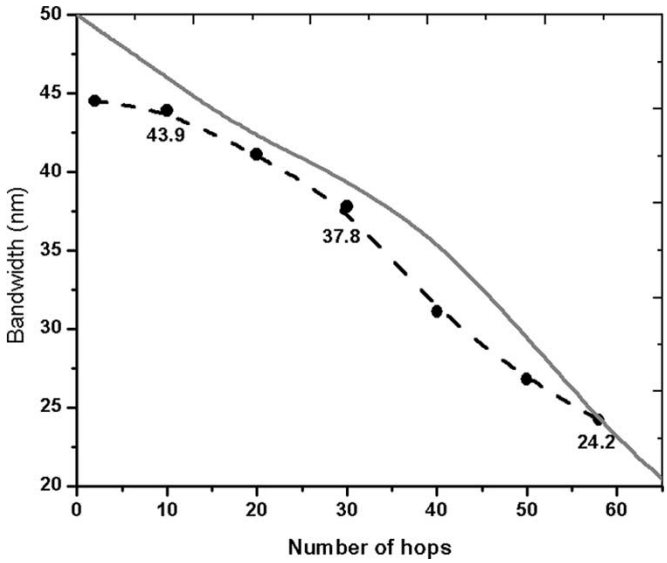

Fig. 10. Functional bandwidth for eight payload channels versus number of hops. The solid gray line represents the simulation results.

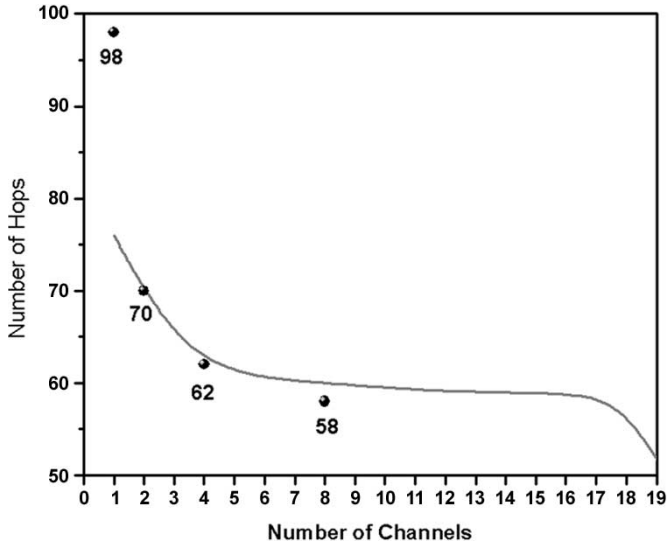

Fig. 11. Experimental results for one, two, four, and eight channels. The maximum number of hops is set for bit error measurement below 10^{-9} for all WDM channels. The solid gray line shows simulation results. Functional bandwidth for eight payload channels versus number of hops. The solid gray line represents the simulation results.

In order for the packet to propagate through multiple SOA-based switching nodes, the total optical average input power of the packet should therefore not exceed the saturation power of the SOA.

The recirculating experimental results also demonstrate that both the total optical power of the packet and the optical power of each channel within the packet play an important role in the scalability limitation of an OPS interconnection network. On one extreme, the OSNRs of the individual channels are reduced by the accumulated ASE noise if lower input powers are launched. Without any optical regeneration at the front-end, the receiver reaches its sensitivity limits [17]. On the other extreme, higher per-channel input powers force the SOAs to operate in the saturation regime, leading to crosstalk between channels and gain compression, both of which strongly degrade the BER performance [39]. The input power of each channel is therefore balanced between these two extremes in order to optimize the network scalability.

In addition to the power restrictions, the wavelength dependences of the SOA gain profiles and of the passive elements are important considerations particularly for WDM OPS networks. With the SOA gain peak centered at 1465 nm, the shorter wavelength channels experience not only more gain but also more ASE noise. After 60 hops (Fig. 7), the payload data in the four shortest wavelength channels are the first to fall short of the performance baseline metric of a BER of 10^{-9}. Longer wavelength channels experience less gain, so they are slightly pre-emphasized. However, the cascading effect of passive optical components cannot be ignored (Fig. 8). As each channel experiences different amounts of optical power gain and losses, balancing the power becomes difficult to achieve.

B. WDM Payload Considerations

Since both gain and loss profiles of the various passive and active components as well as the ASE noise are all wavelength dependent, the functional bandwidth across which the WDM payload channels can be distributed is investigated. Fig. 10 shows the maximum functional bandwidth for which all eight channels exhibit a BER below 10^{-9} as a function of the number of hops. The functional bandwidth increases from 24.2 to 37.8 nm when the packet propagates through 30 hops instead of 58 hops and nearly doubles to 43.9 nm if the packet propagates through ten hops. The modeling code is modified to iteratively make predictions about the effects of the utilized WDM bandwidth on the number of SOA hops, which closely match experimental results.

Fig. 11 shows the experimentally measured dependence of the number of SOA hops on the number of payload channels in a packet. Eight payload channels are first successfully launched through 58 hops. The four outermost channels are then disabled and the number of hops through which the packet propagates is incremented until a 10^{-9} BER baseline is reached. Reducing the number of channels to four allows an increase of four node hops. Similarly, 70 hops can be obtained for a packet with two payload channels, located at 1552.8 and 1555.5 nm, where the ASE noise of the SOA is low. Finally, a single channel at 1555.5 nm can propagate through 98 hops [40].

Two phenomena within the SOA limit the physical layer scalability with respect to the number of payload channels. First, for a constant total input power, the amount of gain provided by optical amplifiers in the linear regime is fixed and independent of the number of channels being amplified. For equal channel power, fewer channels require fewer carrier excitations and more gain is available before the SOA reaches saturation [16]. The amount of gain per channel will affect the ability to sufficiently amplify each channel for loss compensation. Also, once the SOA reaches the saturation regime, the gain provided to each channel is lower than the gain provided in the linear regime. As the ASE noise rises in the channels and the

gain gradually drops due to saturation of the SOA, the OSNR and signal power degrade.

VI. Conclusion

The recirculating loop demonstration is a practical and efficient methodology to investigate the scalability of the physical layer of a WDM optical packet interconnection network. This analysis differs from previous recirculating loop demonstrations in that the focus is on the SOA as a low-gain switching element in an OPS network instead of a booster amplifier in a long-haul transmission system. The investigations presented study of the number of hops through which a packet propagates in the recirculating loop while maintaining a BER below 10^{-9}, as a metric for physical layer scalability. The number of hops is related to the number of payload channels in a packet and to the functional bandwidth over which those channels are distributed. The maximum number of payload channels is primarily determined by the optical gain saturation of the SOA, whereas the functional bandwidth is primarily limited by the cascaded components' spectrum profiles. It is also important to maintain the power balance of a network, but this is complex to achieve as the total launched optical power of the packet and the optical input power of each channel must be balanced between two extremes. On one extreme, high power causes the SOAs to operate in the saturation regime and leads to nonlinear effects and signal distortion. On the other end, a lower input power per channel will degrade the OSNR and stress the receiver sensitivity. Both situations are manifested in the overall degradation of the BER performance and limit the physical layer scalability. Beyond the fine power balance of each channel, the bandwidth across which multiple WDM payload channels can be distributed is found to be limited. Smaller networks allow for broader functional bandwidth, whereas larger networks have a narrower functional bandwidth as the cascading effect of the network components attenuates the outermost wavelength channels. Finally, the maximum number of payload channels in a packet for a physical network size is related to the gain saturation of the SOAs. Channels are optimally placed in the region of the functional bandwidth where the ASE noise from the SOAs is the lowest and where the OSNR and sensitivity requirements of the optical receiver are maintained.

Experiments and simulations show a BER below 10^{-9} for an eight-channel 10 Gb/s-per-channel payload distributed over a bandwidth of 24.2 nm of the C-band propagating through a total of 58 switching elements. This is sufficient for a heavily loaded 10k $\times$ 10k port interconnection network (see Fig. 3), which is the number of ports currently envisioned for future large-scale OPS networks in applications ranging from data communications and storage to high-performance computing [1].

The results have demonstrated that the physical layer size scalability of an SOA-based optical switching network is determined by several parameters: the packet's launched input power, the power per channel, the number of channels, and the bandwidth across which the channels are distributed. Optical packet interconnection networks such as the data vortex offer the bandwidth and the low latency required in such systems. However, the physical layer scalability is a critical performance parameter as switching elements are required to reliably route and propagate payload packets with low BER.

References

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*. San Francisco, CA: Morgan Kaufmann, 1998.

[2] G. I. Papadimitriou, C. Papzoglou, and A. S. Pomportsis, "Optical switching: Switch fabrics, techniques, and architectures," *J. Lightw. Technol.*, vol. 21, no. 2, pp. 384–405, Feb. 2003.

[3] R. S. Tucker and W. D. Zhong, "Photonic packet switching: An overview," *IEICE Trans. Electron.*, vol. E82-C, no. 2, pp. 202–212, Feb. 1999.

[4] J. Trezza, H. Hamster, J. Iamartino, H. Begheri, and C. DeCusatis, "Parallel optical interconnects for enterprise class server clusters: Needs and technology solutions," *IEEE Commun. Mag.*, vol. 41, no. 2, pp. S36–S42, Feb. 2003.

[5] Q. Yang and K. Bergman, "Performances of the data vortex switch architecture under nonuniform and bursty traffic," *J. Lightw. Technol.*, vol. 20, no. 8, pp. 1242–1247, Aug. 2002.

[6] J. Gripp, M. Duelk, J. E. Simsarian, A. Bhardwaj, P. Bernasconi, O. Laznicka, and M. Zirngibl, "Optical switch fabrics for ultra-high-capacity IP routers," *J. Lightw. Technol.*, vol. 21, no. 11, pp. 2839–2850, Nov. 2003.

[7] S. Banerjee and D. Sarkar, "Hypercube connected rings: A fault-tolerant and scalable architecture for virtual lightwave network topology," in *Proc. Information Communications (INFOCOM)*, Toronto, ON, Canada, 1994, pp. 1236–1243.

[8] M. Y. Jeon, Z. Pan, J. Cao, C. Jing, Y. Bansal, J. Taylor, Z. Wang, V. Akella, K. Okamoto, S. Kamei, J. Pan, and S. J. B. Yoo, "Demonstration of all-optical packet switching routers with optical label swapping and 2R regeneration for scalable optical label switching network applications," *J. Lightw. Technol.*, vol. 21, no. 11, pp. 2723–2726, Nov. 2003.

[9] D. J. Blumenthal, B.-E. Olsson, G. Rossi, T. E. Dimmick, L. Rau, M. Mašanoviæ, O. Lavrova, R. Doshi, O. Jerphagnon, J. E. Bowers, V. Kaman, L. A. Coldren, and J. Barton, "All-optical label swapping networks and technologies," *J. Lightw. Technol.*, vol. 18, no. 12, pp. 2058–2074, Dec. 2000.

[10] B. Meagher, G.-K. Chang, G. Ellinas, Y. M. Lin, W. Xin, T. F. Chen, X. Yang, A. Chowdhury, J. Young, S. J. Yoo, C. Lee, M. Z. Iqbal, T. Robe, H. Dai, Y. J. Chen, and W. I. Way, "Design and implementation of ultra-low latency optical label switching for packet-switched WDM networks," *J. Lightw. Technol.*, vol. 18, no. 12, pp. 1978–1987, Dec. 2000.

[11] W. J. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. San Francisco, CA: Morgan Kaufmann, 2004.

[12] R. Hemenway, R. R. Grzybowski, C. Minkenberg, and R. Luijten, "Optical-packet-switched interconnect for supercomputer applications," *J. Opt. Netw.*, vol. 3, no. 12, pp. 900–913, Nov. 2004.

[13] L. Dittmann, C. Develder, D. Chiaroni, F. Neri, F. Callegati, W. Koerber, A. Stavdas, M. Renaud, A. Rafel, J. Solé-Pareta, W. Cerroni, N. Leligou, L. Dembeck, B. Mortensen, M. Pickavet, N. Le Sauze, M. Mahony, B. Berde, and G. Eilenberger, "The European IST project DAVID: A viable approach toward optical packet switching," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 7, pp. 1026–1040, Sep. 2003.

[14] D. K. Hunter, M. H. M. Nizam, M. C. Chia, I. Andonovic, K. M. Guild, A. Tzanakaki, M. J. O'Mahony, L. D. Bainbridge, M. F. C. Stephens, R. V. Penty, and I. H. White, "WASPNET: A wavelength switched packet network," *IEEE Commun. Mag.*, vol. 37, no. 3, pp. 120–129, Mar. 1999.

[15] I. Armstrong, I. Andonovic, and A. E. Kelly, "Semiconductor optical amplifiers: Performance and applications in optical packet switching," *J. Opt. Netw.*, vol. 3, no. 12, pp. 882–897, Nov. 2004.

[16] M. J. Connelly, *Semiconductor Optical Amplifiers*. London, U.K.: Kluwer, 2002.

[17] G. P. Agrawal, *Fiber-Optic Communication Systems*, 3rd ed. New York: Wiley, 2002.

[18] L. H. Spiekman, J. M. Wiesenfeld, A. H. Gnauck, L. D. Garrett, G. N. van den Hoven, T. van Dongen, M. J. H. Sander-Jochem, and J. J. M. Binsma, "8 $\times$ 10 Gb/s DWDM transmission over 240 km of standard fiber using a cascade of semiconductor optical amplifier," *IEEE Photon. Technol. Lett.*, vol. 12, no. 8, pp. 1082–1084, Aug. 2000.

[19] S.Baruh, M. R. X. de Barros, M.de L. Rocha, M. R. Horiuchi, J. B. Rosolem, R. Arradi, S. M.Rossi, A. Paradisi, M. T. M. Rocco Giraldi, and M. A. G. Martinez, "Experimental demonstration and numerical simulation of an optical recirculating loop operating at 10 Gb/s," in *Proc. SBMO/IEEE MTT-S Int.,* Foz do Iguaçu, Brazil, 2003, pp. 239–243.

[20] Y. Sun, I. T. Lima, H. Jiao, J. Wen, H. Xu, H. Ereifej, G. M. Carter, and C. R. Menyuk, "Study of system performance in a 107-km dispersion-managed recirculating loop due to polarization effect," *IEEE Photon. Technol. Lett.*, vol. 13, no. 9, pp. 966–968, Sep. 2001.
[21] A. Tzanakaki and M. J. O'Mahony, "Analysis of tunable wavelength converters based on cross-gain modulation in semiconductor optical amplifiers operating in the counter propagating mode," *Proc. Inst. Elect. Eng.—Optoelectron.*, vol. 147, no. 1, pp. 49–55, Feb. 2000.
[22] Z. Li, Y. Dong, J. Mo, Y. Wang, and C. Lu, "1050-km WDM transmission of 8 × 10.709 Gb/s DPSK signal using cascaded in-line semiconductor optical amplifier," *IEEE Photon. Technol. Lett.*, vol. 16, no. 7, pp. 1760–1762, Jul. 2004.
[23] J. Yu and P. Jeppesen, "Improvement of cascaded semiconductor optical amplifier gates by using holding light injection," *J. Lightw. Technol.*, vol. 19, no. 5, pp. 614–623, May 2001.
[24] K. Bergman and G. Hughes. (2004, Dec.). Optical Interconnection networks. *J. Opt. Netw.* [Online]. 3(12), pp. 760–836. Available: http://www.osa-jon.org
[25] M. Listanti, V. Eramo, and R. Sabella, "Architectural and technological issues for future optical Internet networks," *IEEE Commun. Mag.*, vol. 38, no. 9, pp. 82–92, Sep. 2000.
[26] D. K. Hunter, M. C. Chia, and I. Andonovic, "Buffering in optical packet switches," *J. Lightw. Technol.*, vol. 16, no. 12, pp. 2081–2094, Dec. 1998.
[27] H. J. S. Dorren, M. T. Hill, Y. Liu, N. Calabretta, A. Srivatsa, F. M. Huijskens, and H. de Waardt, "Optical packet switching and buffering by using all-optical signal processing methods," *J. Lightw. Technol.*, vol. 21, no. 1, pp. 2–12, Jan. 2003.
[28] C. Reed, "Multiple level minimum logic network," U.S. Patent 5 996 020, Nov. 30, 1999.
[29] Q. Yang, K. Bergman, G. D. Hughes, and F. G. Johnson, "WDM packet routing for high-capacity data networks," *J. Lightw. Technol.*, vol. 19, no. 10, pp. 1420–1426, Oct. 2001.
[30] B. A. Small, O. Liboiron-Ladouceur, A. Shacham, J. P. Mack, and K. Bergman, "Demonstration of a complete 12-port terabit capacity optical packet switching fabric," in *Optical Fiber Communication (OFC)*, Anaheim, CA, 2005, Paper OWK1.
[31] B. A. Small, A. Shacham, K. Bergman, K. Athikulwongse, C. Hawkins, and D. S. Wills, "Emulation of realistic network traffic patterns on an eight-node data vortex interconnection network subsystem," *J. Opt. Netw.*, vol. 3, no. 11, pp. 802–809, Nov. 2004.
[32] W. Lu, B. A. Small, J. P. Mack, L. Leng, and K. Bergman, "Optical packet routing and virtual buffering in an eight-node data vortex switching fabric," *IEEE Photon. Technol. Lett.*, vol. 16, no. 8, pp. 1981–1983, Aug. 2004.
[33] L. Gillner, E. Goobar, L. Thylen, and M. Gustavsson, "Semiconductor laser amplifier optimization: An analytical and experimental study," *IEEE J. Quantum Electron.*, vol. 25, no. 8, pp. 1822–1827, Aug. 1989.
[34] F. T. Arecchi and R. Bonifacio, "Theory of optical maser amplifiers," *IEEE J. Quantum Electron.*, vol. QE-1, no. 4, pp. 169–178, Jul. 1965.
[35] B. A. Small, J. N. Kutz, W. Lu, and K. Bergman, "Modeling, simulating, and characterizing the performance of data vortex switching nodes," in *Proc. Lasers and Electro-Optics Society (LEOS)*, Tucson, AZ, 2003, pp. 53–54.
[36] O. Liboiron-Ladouceur, W. Lu, B. A. Small, and K. Bergman, "Physical layer scalability demonstration of a WDM packet interconnection network," in *Proc. Lasers and Electro-Optics Society (LEOS)*, San Juan, PR, 2004, pp. 567–568.
[37] N. A. Olsson, "Lightwave systems with optical amplifiers," *J. Lightw. Technol.*, vol. 7, no. 7, pp. 1071–1082, Jul. 1989.
[38] J. Zhou, J. Chrostowski, and P. Myslinski, "Measurements of very low bit-error rates of optical switches based on semiconductor optical amplifiers," *IEEE Photon. Technol. Lett.*, vol. 9, no. 8, pp. 1131–1133, Aug. 1997.
[39] E. Almstrom, G. P. Larsen, L. Gillner, W. H. Van Berlo, M. Gustavsson, and E. Berglind, "Experimental and analytical evaluation of packaged 4 × 4 InGaAsP/InP semiconductor optical amplifier gate switch matrices for optical networks," *J. Lightw. Technol.*, vol. 14, no. 6, pp. 996–1004, Jun. 1996.
[40] W. Lu, O. Liboiron-Ladouceur, B. A. Small, and K. Bergman, "Cascading switching nodes in data vortex optical packet interconnection network," *Electron. Lett*, vol. 40, no. 14, pp. 895–896, Jul. 2004.

Odile Liboiron-Ladouceur (S'95) was born in Montréal, QC, Canada, in 1976. She received the B.Eng. degree in electrical engineering from McGill University, Montréal, in 1999 and the M.S. degree in electrical engineering from Columbia University, New York, NY, in 2003. She is currently working toward the Ph.D. degree in electrical engineering at Columbia University.

From 1999 to 2000, she was an Applications Engineer in the mass storage business unit at Teradyne Inc. She then joined Texas Instruments in 2000 and spent two years working in the fiber optic business unit as a Design Engineer. Her interests include optical interconnection physical layer analysis and clock and data recovery solutions for optical packet-switching interconnection networks.

Benjamin A. Small (S'98) received the B.S. degree (with honors) and the M.S. degree from the Georgia Institute of Technology, Atlanta, in 2001 and 2002, respectively, and the M.Phil. degree from Columbia University, New York, NY, in 2005, all in electrical engineering. He is currently working toward the Ph.D. degree at Columbia University.

His interests include optoelectronic device physics and modeling, as well as optical packet-switching interconnection network traffic analysis and system-level behavior.

Keren Bergman (S'87–M'93) received the B.S. degree in electrical engineering from Bucknell University, Lewisburg, PA, in 1988, and the M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 1991 and 1994, respectively.

She is an Associate Professor of Electrical Engineering at Columbia University. Her research interests include ultrafast optical signal processing, wavelength-division-multiplexing (WDM) optical networking, and optical packet-switched interconnection networks.

Prof. Bergman is a Fellow of the Optical Society of America (OSA). She currently serves as the Associate Editor for IEEE PHOTONICS TECHNOLOGY LETTERS and for the *OSA Journal of Optical Networking*.