

Wavelength division multiplexing based optical backplane with arrayed waveguide grating passive router

Evangelos Grivas

National Center for Scientific Research
Demokritos
Institute of Microelectronics
Agia Paraskevi
Attiki, Greece

Efstathios D. Kyriakis-Bitzaros

National Center for Scientific Research
Demokritos
Institute of Microelectronics
Agia Paraskevi
Attiki, Greece
and
Technological Educational Institute of Piraeus
Department of Electronics
Egaleo, Greece

George Halkias

Stavros G. Katsafouros

National Center for Scientific Research
Demokritos
Institute of Microelectronics
Agia Paraskevi
Attiki, Greece

Geert Morthier

Pieter Dumon

Roel Baets, MEMBER SPIE

University of Ghent
Department of Information Technology (INTEC)
Ghent, Belgium

Tom Farell

Neil Ryan

Intune Technologies
Dublin, Ireland

Iain McKenzie, MEMBER SPIE

Errico Armadillo

European Space Agency
European Space Research and Technology
Centre (ESA/ESTEC)
Nordwijk, The Netherlands

1 Introduction

Wavelength division multiplexing (WDM) is a well-known and widely applied technique in long-haul telecommunica-

Abstract. A wavelength division multiplexing (WDM)-based optical backplane architecture is introduced. It is a tunable transmitter fixed receiver (TT-FR) architecture incorporating an $N \times N$ arrayed waveguide grating (AWG) element for passive data routing between the nodes, which in conjunction with star couplers, offers both unicast and multicast capabilities. The data and control plane of the network are implemented on a high-speed field programmable gate array (FPGA), and a four-node demonstrator is built up. Three different types of AWG routing elements implemented in different technologies are employed, and bit error rate (BER) versus incident power on the receiver measurements are presented for a data rate of 10 Gbps per link. A total switching time as low as 500 ns is achieved, permitting packet switching operation with more than 95% efficiency when the packet length is greater than 10 kbytes. © 2008 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.2842380]

Subject terms: optical backplane; wavelength division multiplexing; passive optical routing; optical interconnects.

Paper 070182R received Mar. 1, 2007; revised manuscript received Sep. 7, 2007; accepted for publication Sep. 12, 2007; published online Feb. 26, 2008.

tion networks¹⁻³ and local area computer networks,⁴⁻⁶ which offers increased bandwidth through a single optical fiber. However, the advantages of WDM are not extensively investigated for board-to-board or even chip-to-chip interconnects. The application of networking concepts to optoelectronic multiprocessor architecture design offer advan-

tages to the system designer, since already tested solutions to the problems occurring in the development of complex hybrid architectures and associated protocols may be transferred from one domain to the other.⁷ In Ref. 8, a WDM-based optical backplane for high-speed IP routers is proposed using passive routing of data packets between network interface cards. A prototype system working at a data rate of 40 Gbps is presented using a PC-based controller to guarantee collision-free operation. In Refs. 9 and 10, the use of WDM principles is proposed to overcome the bandwidth restrictions in the memory-to-processor and the processor-to-processor interconnection requirements of a hierarchical multiprocessor system, while a time-deterministic WDM star network for application in parallel signal processing of radar systems is proposed in Ref. 11.

In this work we present a WDM optical backplane architecture oriented for application to on-board satellite signal and data processors. The requirement for on-board data and signal handling is expected to increase dramatically in the near future by the introduction of high-data-rate sensors, such as the next generation of synthetic aperture radars (SARs), multispectral imagers, and high-speed high-resolution cameras requiring vast amounts of memory and on-board signal processing. On-board processors perform tasks such as data selection and compression before subsequent downlinking to decrease the required communication bandwidth. A four-node demonstrator has been set up using tunable transmitters and fixed receivers (sensitive to the entire wavelength range of interest) and high-speed field programmable gate arrays (FPGAs) implementing a hybrid centralized-distributed control mechanism. Three different arranged waveguide grating (AWG) elements,¹² which exhibit different characteristics in terms of multicast capability, degree of nonblocking, and available bandwidth, have been used. Bit error rate (BER) measurements versus power incident on the receiver are initially performed for static backplane operation. Testing under dynamic reconfiguration conditions is carried out subsequently, and the capability of packet switching operation is demonstrated.

The rest of this work is organized as follows. In Sec. 2 the optical backplane architecture is introduced and the details of the implementation of a four-node prototype system are discussed. In Sec. 3 the BER measurements for different system configurations are presented. Finally, the conclusions are drawn in Sec. 4.

2 Backplane Architecture and Implementation

2.1 Architecture

Depending on the tunability properties of the transmitters and the receivers, WDM network architectures can be classified into four categories¹: tunable transmitter tunable receiver (TT-TR), tunable transmitter fixed receiver (TT-FR), fixed transmitter tunable receiver (FT-TR), and fixed transmitter fixed receiver (FT-FR).

The proposed backplane is a TT-FR architecture, meaning that any transmitting node can be tuned to send data to any of the permissible wavelengths, while the receiver accepts incoming data in all wavelengths. The basic architecture of the WDM backplane is depicted in Fig. 1(a). It mainly consists of two parts: the wavelength router and the processing nodes communicating through the backplane.

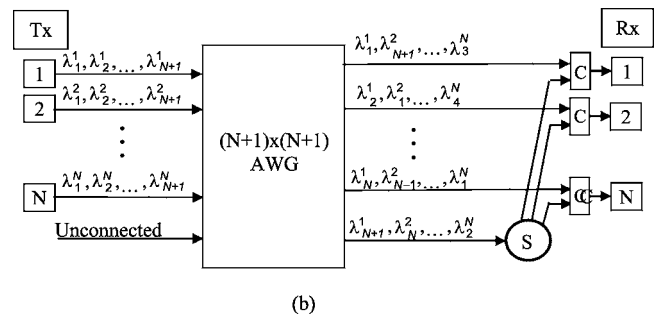
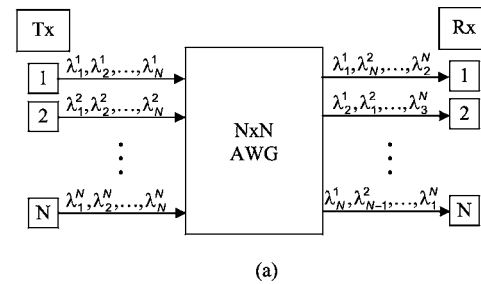


Fig. 1 Architecture of AWG-based optical backplane: (a) unicast only, and (b) broadcast added with extra AWG inputs, splitter (S), and combiners (C).

For illustration clarity, the transmitter and the receiver part of each node are drawn separately. Moreover, referring to Fig. 1, the subscript on λ characterizes the different wavelengths, while the superscript denotes the originating node.

The router is an arrayed waveguide grating (AWG) (also termed waveguide grating router, WGR), which is a passive component within principle N input ports and N output ports.¹² It is different from a star coupler by its wavelength dependence: different wavelengths at a certain input port are routed to different output ports, as illustrated in Fig. 1(a). Thus, from any input port one can address any output port by selecting the appropriate wavelength.

By using an $N \times N$ AWG with number of I/Os larger than the number of nodes of the network, or an $N \times M$ AWG with $M > N$ and adding some extra passive components (splitters and combiners), broadcast and/or multicast features can also be implemented. In Fig. 1(b) the implementation of broadcasting in an N -node network using an $(N+1) \times (N+1)$ AWG, a 1: N splitter, and N 2:1 combiners is depicted.

The passive wavelength router transparently implements the network connectivity on the physical level together with the laser circuitry. By using this combination of tunable transmitters, fixed receivers, and fixed communication paths in the router, responsibility for adding more flexibility in the network is put on higher network levels. Although the $N \times N$ AWG operating with a unique set of N wavelengths is a nonblocking routing element,¹² collisions may arise on the overall backplane when two or more nodes are transmitting simultaneously to a single node. Thus, a simple arbitration mechanism or a complete medium access control (MAC) protocol should be used to guarantee the collision-free operation of the backplane.

Table 1 Basic characteristics of the three AWGs.

AWG type	Number of I/Os	Size (mm ²)	Insertion loss (dB)	Cross talk (dB)	Temperature dependence	Polarization dependence
SOS	4×4	3000	7 to 9	-25	low	low
SOI	4×4	2	12 to 14	-12	high	high
Polymer	8×8	1000	<12	<-20	low	high

2.2 Implementation

A four-node prototype system implementing both architectures shown in Fig. 1 has been built up. Three different AWGs have been used, two of them were 4×4, and the third one was 8×8 and, combined with an external star coupler and four combiners, was used to offer broadcast capabilities in the system. Moreover, to enable fast switching experiments and reduce prototype implementation cost, two of the nodes had conventional transponders with sending and receiving capability, while the other two nodes possessed burst-mode receivers only.

2.2.1 Wavelength router

The three AWGs used and tested are a commercially available 4×4 AWG in silica-on-silicon (SoS) technology, a research prototype 4×4 AWG in silicon-on-insulator (SOI) technology,¹³ and another prototype of 8×8 AWG in polymer technology. Some basic properties of the AWGs have been measured before their application to the backplane system and the results are summarized in Table 1.

Comparing the characteristics of the three AWGs used, it should be mentioned that while the SoS component has the largest footprint (~30 cm²), it presents great performance since the minimum insertion loss is quite low (7 dB) and the cross talk generated between communication channels is lower than -25 dB. This means that the technology is scalable to a larger number of communication channels. Its polarization dependency and temperature dependency are very low, making the technology easily employable. On the other hand, the SOI component is significantly smaller with a total die size of just 2 mm², but its minimum insertion loss is 12.5 dB, and the cross talk induced between channels ranges from -12 to -14 dB. The SOI component is very polarization and temperature sensitive and was therefore used with polarization-controlled inputs and a thermal stabilization system. Finally, the polymer routing component has properties somewhere in-between those of the SoS and SOI components; its footprint is about one third of that of the SoS while supporting twice as much channels, but it is still much larger than the SOI component. Its minimum insertion loss is 12 dB, quite close to that of the SOI component and worse than the SoS router. It is quite polarization dependent but not very temperature dependent. The enhanced backplane functionality (broadcast and/or multicast features) that can be realized using the polymer AWG combined with external components could be integrated on the SOI chip without noticeably growing the die size.

2.2.2 Transponders

The transponder modules consist of a tunable transmitter and a fixed receiver capable of sending and receiving 10 Gbps data streams while they communicate electrically with the data and control module described in the next section through a 300-pin multi-source agreement (MSA) compatible interface. The transponders used employ a common off-the-shelf (COTS) laserless transponder driven by external fast wavelength tunable laser sources. The lasers operate in the C band according to the ITU-T G.692 standard, i.e., there are 80 wavelength channels from 1528 to 1563 nm with 50-GHz spacing between each other. When switching, the output of the laser is blanked to avoid disrupting other communication channels, and when blanking is removed, the laser is actively locked to the correct wavelength with the use of an integrated Fabry Perot wavelength locker to provide a reference signal. For all wavelength transitions over the C band, i.e., from each unique wavelength to every other wavelength, the maximum switching time is <200 ns, as shown in the histogram of Fig. 2. The wavelength is specified to be within ±2.5 GHz of the expected value for such transitions cases. As the demonstrator uses only a small subset of these wavelengths, the typical wavelength switching time is less than 100 ns.

The receivers are based on PIN photodiodes. Two of them are part of the transponders incorporating an optical amplifier for dynamic range extension and presenting a lock-up time of the clock and data recovery (CDR) circuit on the order of several microseconds. The other two are in-house-developed burst-mode receivers with CDR lock-up time of less than one microsecond. The exact

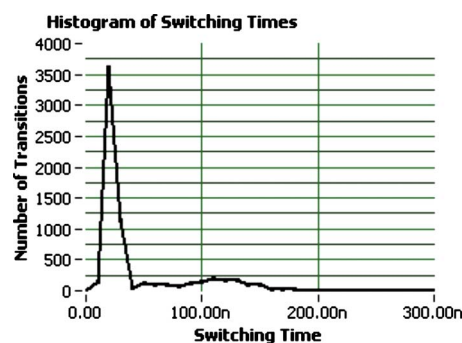


Fig. 2 Histogram of switching times for all possible wavelength transitions of the tunable laser.




Fig. 3 Block diagram of the node hardware.

lock-up time depends on the illumination conditions, and as expected is greater when the receiver starts to operate after a period of darkness.

2.2.3 Data generation and control units

The control and data planes of the backplane have been implemented using the Stratix-II® high-speed FPGAs of Altera (San Jose, California). The block diagram of the system is shown in Fig. 3. Concerning the data plane, every node generates outgoing data packets, processes incoming data, and performs bit error rate (BER) measurements. The data packet is comprised of a variable length receiver lock-up pattern, a fixed start-of-data pattern (SoDP), and a $2^{23} - 1$ pseudorandom bit sequence. To calculate the BER of the incoming data stream as soon as the receiver recognizes the SoDP, a pseudorandom bit sequence generator identical to that of the transmitter is initiated, and received data are compared to expected data. Every discrepancy between the two data streams is registered to the error counter.

The internal data path is implemented using a 64-bit architecture operating at a core frequency of 155.52 MHz to produce 16 bit streams at 622.08 Mbps, which are then multiplexed within the transponder at the serial data rate of 9.953 Gbps of the optical signal. The 300-pin MSA interface is used to interface the FPGA board with the optical transponders.

A preallocation MAC protocol has been implemented to enable collision-free system operation. A hybrid centralized/distributed control mechanism has been used, where all nodes contain predefined routing tables and communication schemes, and a master node is responsible for triggering events and controlling synchronization between

the nodes. The various configurations required to test the backplane have been stored in the FPGA memory. Specifically, 51 configurations for static measurements and three multistep test suites for functional and optical burst-mode switching tests have been implemented. Moreover, the controller generates the signals required to select the appropriate wavelength address on the tunable laser module through a proprietary fast wavelength assignment interface. Finally, the FPGA boards are also connected through a JTAG interface to a PC for programming the FPGA devices and for data retrieval during the tests.

A photograph of the four-node setup of the demonstrator is shown in Fig. 4. This involves the four FPGA controller boards, the two laserless transponders, two burst-mode receivers, and two tunable laser sources. Thus, a four-node system with two full duplex (transmit and receive) nodes and two receive-only nodes is implemented. All nodes operate at a full data rate of 10 Gb/s.




Fig. 4 Photograph of the four-node demonstrator.




Fig. 5 10-Gbps optical eye diagram with OC192 SONET mask for optical loopback test of node 0.

The use of the off-the-shelf components for the build up of the demonstrator, although it gives us the capability to estimate the overall system performance and evaluate the properties of the different AWGs used, imposes increased requirements in terms of electrical I/Os for the processing units. However, in a real-life application, the functions of the transponder modules (clock and data recovery, bit-stream encoding, serializer-deserializer, etc.) should be integrated on a single chip. This can be easily done exploiting the capabilities of the contemporary nanometer-scale complementary metal oxide semiconductor (CMOS) technologies (there are already commercially available FPGAs incorporating multiple multigigabit transceivers). Furthermore, the integration of the active optical devices with the electronic circuits can be performed either at the chip level (using hybrid integration techniques¹⁴ or silicon photonics¹⁵) or at the board level,¹⁶ resulting in compact and efficient solutions that overcome the drawbacks of the proposed implementation. The development of complete monolithic photonic circuits¹⁵ may lead to even greater degrees of integration.

3 Test Results

The proper operation of every node has been verified, initially, using electrical and optical loopback connections. No errors have been recorded for more than 10 h of continuous operation leading to a BER $< 10^{-15}$. A typical optical signal eye diagram at 10 Gbps, which reveals the large margin of error-free operation, is shown in Fig. 5.

Evaluation of the overall system performance and comparison of the three routing elements were performed subsequently. External attenuation of the optical signal was used to obtain BER versus received power curves. The first set of measurements aimed to get the single channel BER figure by setting up a single point-to-point connection through the routing element. The results are summarized in Fig. 6. The SoS component presents the smaller variation between different channels (< 0.2 dB), while the polymer and the SOI AWGs are more sensitive to wavelength change. The variation observed for the SOI AWG is partially due to the fact that the tunable laser output and the AWG input wavelength are not exactly matched in all

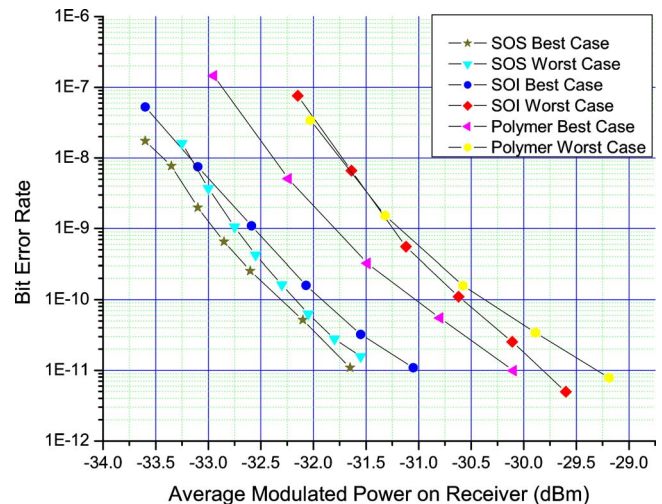


Fig. 6 Back-to-back BER measurements for the SOS, SOI, and polymer AWGs.

cases. Unicast and broadcast BER measurements performed using the polymer AWG showed a 0.5-dB power variation between all four channels for broadcast operation, which is mainly due to insertion loss variations of the external couplers used to implement the broadcasting scheme.

The next set of experiments intended to determine the influence of incoherent and coherent cross talk, and the power penalty introduced by single and multiple channel cross talk. Initially, power penalty due to single source cross talk was measured. The cross talk level (CL) was being adjusted, when necessary, by attenuating the optical power of the appropriate transmitter. The power penalty for incoherent cross talk was found to be 0.05 dB for the SoS AWG at a CL of 25 dB, and 0.2 dB for the polymer device at a CL of 22.5 dB. For the SOI AWG, the power penalty varied from 0.1 dB at a CL of 19 dB to 1 dB at a CL of 11.1 dB. The power penalty due to coherent cross talk presented greater variations, as expected. For the SoS device, it was in the range of 0.05 to 0.075 dB (for different optical paths) at a CL of 25 dB, and increased to 0.2 to 0.3 dB as the CL decreased to 13.4 dB with attenuation of the optical signal. The SOI AWG gave a 2-dB power penalty at a CL of 14.1 dB that reduced to 0.2 dB as the CL increased to 20.1 dB by attenuating the signal causing the cross talk. The polymer device experienced the greater sensitivity to coherent cross talk as the power penalty was 1.25 dB at a CL of 22.5 dB, and further increased to 3 dB at a CL of 20 dB. All power penalties were measured at a BER of 10^{-10} . Finally, an external modulated tunable source was used with the SoS AWG as a third transmitting node, since as mentioned earlier only two transponders were available. The cumulative power penalty due to the two cross talk sources increased to 0.45 dB at a CL of 13.4 dB (Fig. 7), which implies the effect is additive.

The tests reported up to now are static, since the back-plane configuration does not change during the test. A set of dynamic tests has also been performed using the burst-mode receivers of nodes 2 and 3 to examine packet switching capabilities of the system. Two scenarios have been examined: always-on and off-to-on switching. In the




Fig. 7 BER measurements under the presence of coherent cross talk (SOS).

always-on case, nodes 0 and 1 send data to nodes 2 and 3, respectively, for a specific time interval, and after that a switching occurs, so that node 0 sends to 3 and 1 to 2. In the off-to-on case, the receiver under test gets data from a transmitting node (i.e., node 0), gets no data for some time, and finally receives data from another node (i.e., node 1). This represents the worst case, since the receiver locks on the incoming data after a period of darkness. In both cases, constant cross talk is always present, as both transmitters are always on.

The optical packet switching (OPS) performance of the burst-mode receivers was investigated to reveal how the full demonstrator performs under packet switched loading, the situation that would exist in a real application. Because the tunable laser modules employed are guaranteed to switch wavelength in <200 ns, the main system limiting functional block, in terms of packet switching capability, is the task of clock and data recovery (CDR) at the burst-mode receiver. The limiting factor on this task (without distributing a clock) is how adaptive the PLL governing the CDR is at synchronizing with the incoming data stream. While the receivers were built with standard MSA compatible components, the CDR status flag signal response is limited to the range of about 40 to 50 μ s. However, actual clock recovery occurs faster than this, commensurate with the frequency and jitter tolerances required, according to the communications protocol and payload type employed. To accurately verify the true lock-in time, bit errors at the receiver are monitored while decreasing the preamble length. The minimum achievable switching time is thus indicated by the minimum preamble length that leads to a BER curve matching that of the respective static test. As depicted in Fig. 8, BER measurements for the off-to-on switching case have shown no degradation at backplane operation using a preamble length of 500 ns and a packet length of 4 sec. Thus, fast switching in the range of 600 ns (assuming laser switching time <100 ns) can be achieved, which for the data rate of 10 Gbps leads to an efficiency of 95% when the packet size is greater than 10 kbytes.

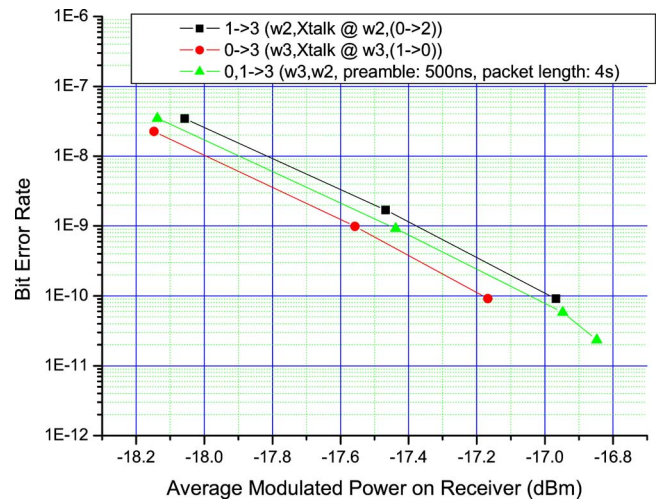


Fig. 8 BER measurements under burst-mode switching operation.

4 Conclusions

A four-node wavelength-routed multigigabit optical backplane, based on a passive AWG router and tunable transponders, is demonstrated. Three passive routers implemented in silica-on-silicon (SoS), in silicon-on-insulator (SOI), and in polymer technologies, respectively, are successfully tested. A preallocation MAC protocol is adopted and realized using high-speed FPGAs. BER measurements versus optical power for static and dynamic backplane configuration using 10-Gbps data streams are performed and optical burst-mode switching is successfully demonstrated. The results reveal the scalability of the architecture to a large number of ports with 40 Gb/s data rate per channel and switching times less than 50 ns.

Acknowledgment

This work was supported by ESA/ESTEC under contract number 17884/03/NL/HE, "Multigigabit Optical Backplane Interconnections."

References

1. M. Maier, *Metropolitan Area WDM Networks: An AWG Based Approach*, Kluwer Academic Publishers, New York (2004).
2. C. Fan, M. Maier, and M. Reisslein, "The AWG|PSC network: a performance-enhanced single-hop WDM network with heterogeneous protection," *J. Lightwave Technol.* **22**(5), 1242–1262 (2004).
3. D. Wonglumsom, I. M. White, K. Shrikhande, M. S. Rogge, S. M. Gemelos, F.-T. An, Y. Fukashiro, M. Avenarius, and L. G. Kazovsky, "Experimental demonstration of an access point for HORNET-A packet-over-WDM multiple-access MAN," *J. Lightwave Technol.* **18**(12), 1709–1717 (2000).
4. B. Mukherjee, "WDM-based local lightwave networks part I: single-hop systems," *IEEE Trans. Neural Netw.* **6**(3), 12–27 (1992).
5. A. V. Krishnamoorthy, J. E. Ford, F. E. Kiamilev, R. G. Rozier, S. Hunsche, K. W. Goossen, B. Tseng, J. A. Walker, J. E. Cunningham, W. Y. Jan, and M. C. Nuss, "The AMOEBA switch: an optoelectronic switch for multiprocessor networking using dense-WDM," *IEEE J. Sel. Top. Quantum Electron.* **5**(2), 261–275 (1999).
6. A. Pattavina, "Architectures and performance of optical packet switching nodes for IP networks," *J. Lightwave Technol.* **23**(3), 1023–1032 (2005).
7. J. Rorie, P. J. Marchand, J. Ekman, F. E. Kiamilev, and S. C. Esener, "Application of networking concepts to optoelectronic multiprocessor architectures," *IEEE J. Sel. Top. Quantum Electron.* **5**(2), 353–359 (1999).
8. J. Gripp, M. Duell, J. E. Simsarian, A. Bhardwaj, P. Bernasconi, O. Laznicka, and M. Zirngibl, "Optical switch fabrics for ultra-high-capacity IP routers," *J. Lightwave Technol.* **21**(11), 2839–2850 (2003).

- (2003).
9. K. M. Sivalingam and P. W. Dowd, "A multilevel WDM access protocol for an optically interconnected multiprocessor system," *J. Lightwave Technol.* **13**(11), 2152–2167 (1995).
 10. P. W. Dowd, K. Bogineni, K. A. Aly, and J. A. Perreault, "Hierarchical scalable photonic architectures for high-performance processor interconnection," *IEEE Trans. Comput.* **42**(9), 1105–1120 (1993).
 11. M. Jonsson, A. Ahlander, M. Taveniku, and B. Svensson, "Time-deterministic WDM star network for massively parallel computing in radar systems," *Proce. MPP01* pp. 85–93 (1996).
 12. P. Bernasconi, C. Doerr, C. Dragone, M. Cappuzzo, E. Laskowski, and A. Paunescu, "Large $N \times N$ waveguide grating routers," *J. Lightwave Technol.* **18**(7), 985–991 (2000).
 13. P. Dumon, W. Bogaerts, D. Van Thourhout, G. Morthier, R. Baets, P. Jaenen, S. Beckx, J. Wouters, T. Farrell, N. Ryan, E. Grivas, E. Kyriakis-Bitaros, G. Halkias, and I. McKenzie, "A nanophotonic 4×4 wavelength router in silicon-on-insulator," 2006 OFC/NFOEC Meeting, Anaheim, CA, paper OTHO5.
 14. A. V. Mule, R. A. Villalaz, P. J. Joseph, A. Naeemi, P. A. Kohl, T. K. Gaylord, and J. D. Meindl, "Polyolithic integration of electrical and optical interconnect technologies for gigascale fiber-to-the-chip communication," *IEEE Trans. Adv. Packag.* **28**(3), 421–433 (2005).
 15. R. Soref, "The past, present, and future of silicon photonics," *IEEE J. Sel. Top. Quantum Electron.* **12**(6), 1678–1687 (2006).
 16. L. Schares et al., "Terabus: terabit/second-class card-level optical interconnect technologies," *IEEE J. Sel. Top. Quantum Electron.* **12**(5), 1032–1044 (2006).

Evangelos Grivas received the BS degree in physics from the University of Thessaloniki, Greece, in 2002, and the MSc in microelectronics from the University of Athens, Greece, in 2004. He is currently working toward the PhD degree at the Optical Communications Laboratory, Department of Informatics and Telecommunications, University of Athens. His research interests include optical transmission over multimode fibers, dispersion compensation by electronic means, and high-speed electronics supporting optical communications.

Efstathios D. Kyriakis-Bitaros received his BSc degree in physics from the University of Patras, Greece, in 1987, and the PhD degree in electrical engineering from the University of Patras, Greece, in 1994. In 1996 he joined the Institute of Microelectronics, National Center for Scientific Research (NCSR) "Demokritos," Athens, Greece, where he is currently an associate researcher. Since 2003 he has been an assistant professor with the Department of Electronics, Technological Educational Institute (TEI) of Piraeus. His current research interests include the development of efficient circuit models for high-speed light emitting devices, the implementation of short-range optical interconnection schemes, and the applications of low-power CMOS circuits. He is the author or coauthor of two book chapters, and several publications in scientific international journals and conference proceedings. He has also served as a reviewer in international journals and conferences. He is a member of the IEEE LEOS Society.

Stavros G. Katsafouros holds a BSc and MSc in electrical engineering, both from New York University. He has well over 35 years of experience in microelectronics in staff and managerial capacities. He was very instrumental in the creation of the Institute of Microelectronics at the National Center for Scientific Research "Demokritos," Athens, Greece, first as an advisor in microelectronics to the Minister of Research and Technology during 1983 to 1984, and subsequently as staff member, associate director, and acting director of the institute from 1984 to the present. He has served, as participant and technically responsible scientist, in numerous European Union and nationally funded research and development programs. His most recent interests are in the area of mixed VLSI design for general and optoelectronic specific applications.

Geert Morthier received the degree in electrical engineering and the PhD degree from the University of Gent in 1987 and 1991, respectively. Since 1991 he has been a member of the permanent staff of Interuniversity Microelectronics Center (IMEC), Belgium. His main interests are in the modeling and characterization of optoelectronic components. He has authored or coauthored more than 150 papers in the field and holds several patents. He is also one of the two authors of the *Handbook of Distributed Feedback Laser* (Artech House, 1997). In 2001 he was appointed part-time professor at Ghent University, where he teaches courses on optical fiber communication and lasers. He serves on the program committee of several international conferences.

Pieter Dumon obtained the MS in electrical engineering degree from Ghent University in 2002. He obtained the Doctorate degree in electrical engineering in 2007, also from Ghent University, with the PhD thesis "Ultra-compact integrated optical filters in silicon-on-insulator by means of wafer-scale technology." He now coordinates the ePIXnet silicon photonics platform, which gives access to the CMOS facilities at IMEC (B) and CEA-LETI (F) for wafer-scale research and prototyping of silicon photonic components. He is a member of the IEEE and the OSA. He has authored and coauthored 22 publications in international journals and more than 50 conference publications.

Neil Ryan obtained his degree in electronic and computer engineering (BAI, honors) from Trinity College Dublin in 1998. He subsequently obtained his Masters in electronic engineering (MEngSc) from University College Dublin, based on research into the application of digital signal processing in medical imaging, specifically ophthalmology. In 2001, he started work with Intune Networks, an optoelectronics company specializing in the delivery of tunable solutions to the telecommunications and gas sensing industries. He specializes in the characterization and control of fast switching, multisection tunable lasers. He has previously participated in both ESA and European IST projects based around tunable lasers. He has authored or coauthored 11 publications and conference proceedings, and two granted or pending patents.

Iain McKenzie received the BEng degree in electronic and electrical engineering and the MPhil degree in optoelectronics from the University of Strathclyde, Glasgow, in 1991 and 1993, respectively. Between 1995 and 2002 he worked as a research scientist first for The Centre for Advance Materials Technology and subsequently for The CoE in Structural Mechanics at Monash University, Melbourne. Since 2002, he has been working as a contractor with the European Space Agency, in the optoelectronic section of The Directorate of Technical and Quality Management. His research interests include optical communications, optical fiber sensors, microphotonics, and optoelectronic packaging for harsh environments.

Biographies of the other authors not available.