

Optical Interconnections within Modern High-performance Computing Systems

Rick Lytel, Howard Davidson, Nyles Nettleton, and Theresa Sze

Optical Interconnections within Modern High-performance Computing Systems

Rick Lytel, Howard Davidson, Nyles Nettleton, and Theresa Sze

SMLI TR-2000-84

May 2000

Abstract:

Optical technologies are ubiquitous in telecommunications networks and systems, providing multiple wavelength channels of transport at 2.5-10 Gbps data rates over single fiber-optic cables. Market pressures continue to drive the number of wavelength channels per fiber and the data rate per channel. This trend will continue for many years to come as e-commerce grows and enterprises demand higher and reliable bandwidth over long distances. E-commerce, in turn, is driving the growth curves for single processor and multi-processor performance in database transaction and web-based servers. Ironically, the insatiable taste for enterprise network bandwidth, which has driven up the volume and pushed down the price of optical components for telecommunications, is simultaneously stressing computer system bandwidth—increasing the need for new interconnection schemes—and providing for the first time commercial opportunities for optical components in computer systems.

This paper will center primarily on the use of optical interconnects within commercial digital computing systems, particularly workstations and servers, and will address mainly board-to-board interconnects within a single cabinet or box. We feel this is the most likely utilization of optics in commercial computer systems for the next decade. We will also provide a practical analysis of inter- and intra-chip optical interconnects and the difficulties they face in real systems.

Keywords—vertical cavity laser, microprocessor, optical interconnect, free-space interconnect



M/S MTV29-01
901 San Antonio Road
Palo Alto, CA 94303-4900

email addresses:

rick.lytel@eng.sun.com
howard.davidson@eng.sun.com
nyles.nettleton@eng.sun.com
theresa.sze@eng.sun.com

© 2000 Sun Microsystems, Inc. All rights reserved. The SML Technical Report Series is published by Sun Microsystems Laboratories, of Sun Microsystems, Inc. Printed in U.S.A.

Unlimited copying without fee is permitted provided that the copies are not made nor distributed for direct commercial advantage, and credit to the source is given. Otherwise, no part of this work covered by copyright hereon may be reproduced in any form or by any means graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an information retrieval system, without the prior written permission of the copyright owner.

TRADEMARKS

Sun, Sun Microsystems, and the Sun logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

For information regarding the SML Technical Report Series, contact Jeanie Treichel, Editor-in-Chief <jeanie.treichel@eng.sun.com>. Reports in this series are also available online on the World Wide Web at <http://www.sun.com/research/>.

Optical Interconnections within Modern High-performance Computing Systems

Rick Lytel, Howard Davidson, Nyles Nettleton and Theresa Sze

Sun Microsystems Laboratories
901 San Antonio Rd.,
Palo Alto, CA 94303

I. Introduction

Modern fiber-optics has provided the network system designer with the tools to provide bandwidth at will and over any distance on the planet [1]. Fiber is the technology of choice for any distance greater than a few tens of meters in LANs, WANs, and trunk applications. The demand for bandwidth in the internet is doubling every 100 days, and e-commerce is still a nascent business. This bodes well for telecommunications equipment and service providers, as it has a direct impact on their net income. At \$0.05-\$0.10 a minute for a long-distance call, an OC(192) SONET link with many wavelength channels carries on order of \$1M of revenue per hour. This implies that service providers are willing to pay generously for costly transmission and termination equipment, as well as amplifiers, repeaters, and switches. In telecommunications, the revenue is in the traffic; the equipment infrastructure and its expensive bandwidth upgrades is a given.

The exact opposite situation prevails in computer systems. The impact of increasing bandwidth demand from the network drives the I/O bandwidth of servers and demands increasing performance from computer systems, specifically processors and memory. Equipment suppliers cannot charge enterprise customers by the bit; instead, the business model is based on the relative computing performance of new generations of machines with respect to prior generations of machines. Internet service providers (ISPs) are concerned mainly with equipment cost and I/O bandwidth. The former is largely determined by Moore's law, while the latter is fixed by the LAN providing the network-computer interface. The computer is assumed to have ever-increasing performance and reliability, while simultaneously costing less as generations evolve.

This difference in business models has largely limited the use of optical components in modern computing machines because computer system designers cannot afford to pay for optical devices that are ubiquitous in telecommunications. But as optical components work their way down the food chain toward higher speed local area networks, component costs drop (and reliability increases) to levels where computer systems designers can start to think about using optical links in their machines. As we approach the millennium, we find that certain optical solutions are now desirable and cost-effective in computer cluster configurations. It may be only a matter of years before optics begins to invade the compute environment. Where will this invasion take place, and why? We pose some answers to these important questions and provide our best guesses for how this invasion will evolve.

II. Optics Outside the Box

The LAN industry has developed standards for links connecting computers to the network and to storage devices [2]. Ethernet has evolved from 10 Mbps electrical specifications through Gigabit Ethernet (1 Gbps) which specifies both electrical and optical interfaces. Though the latter specification has only recently been completed, the 10-Gigabit Ethernet working group is already discussing a standard. A plethora of vendors are now looking at providing an array of solutions, including single line 10 Gbps serializer/deserializer (SerDes), four parallel 2.5 Gbps channels, and even a 'coarse' wavelength division multiplexed, four channel, 2.5 Gbps per channel solution. All such solutions are based upon 850 nm, short-wavelength vertical cavity surface emitting lasers (VCSELs), since these devices are in mass production (millions per year), reliable, easily coupled to fiber, and easy to drive [3]. Simultaneously, the Fibre Channel standard has emerged as the premier computer storage interconnect solution, and all major computer equipment suppliers now offer systems with Fibre Channel to RAIDs (redundant arrays of inexpensive disks). Though fundamentally different in the traffic they handle, both Ethernet and Fibre channel have provided an entry path for optical components into computer systems. And at this point, there is no going back.

Moving inward from the network to the computer, parallel ribbon fiber links for point-to-point interconnection between cabinets have emerged in the past two years. Devices based on the mechanical transfer (MT) ferrule are available at 1.25 Gb/s per fiber [4] and 2.5 Gb/s per fiber [5].

These links are used to extend the cache coherence domain between cabinets, and are generally used in a “zero-protocol” mode. Early adoption of these devices may have been more strongly driven by the need to conserve precious panel space on adapter cards than any technical or cost superiority over existing shielded twisted pair solutions.

Given the emergence of link standards for short wavelength optical solutions in the LAN and storage applications, it was only a matter of time before computer systems designers would attempt to incorporate such technology into their computer clusters, mainly as point-to-point optical links between distributed CPU systems and electronic switches. Recently, two emerging I/O computer interconnect standards, NGIO (next generation IO, driven mainly by Intel and Sun) and Future IO (driven mainly by IBM™ and HP) have merged into the so-called Infiniband standard in an attempt to create a very high market demand for optical (and very high performance electrical) transceivers and cable. Bandwidths for fat versions of such links approach 10s of Gbps, far short of what is used in telecommunications transport, but significantly faster than what has been available to computer system designers. Systems designers can now provide their customers with multiple boxes of clustered or shared-memory CPUs, or even collections of disparate CPUs communicating through a common interconnect standard. The rapid implementation of these optical interconnect standards implies that short-wavelength optical components have broken into the computer interconnect applications and bootstrapped themselves into high-volume, low-cost devices. The exponential growth in the demand for bandwidth in the network has driven the bandwidth requirements for LANs which have, in turn, driven the requirement for relatively short (100m) optical cabling and low-cost I/O optical modules at the machine I/O interface and now, with Infiniband, at the CPU interface. Bandwidth provided the need, but standards have provided the means for production of high-performance, low-cost optical components in computers.

III. Optics Inside the Box

The conclusion of the last section suggests two obvious ways to incorporate optics into the computer system itself, that is, in the box. These are: (1) find an application where the increasing bandwidth demand of the link makes optics the preferred engineering choice, and (2) make the interconnect a standard to drive up volume, increase the supplier base, and lower costs. Since the

component industry has already provided lower cost, short-wavelength multimode lasers, links, and receivers running at a few Gbps each, we focus here on this technology only. Though other solutions exist—long-wavelength lasers, DWDM solutions, all-optical switches—they are very expensive and are not able to lead the invasion of optics into the compute box, though they may eventually comprise part of the total component market.

To begin, we examine the interior of a typical multiprocessor compute box and look for the places where external bandwidth demands impact system design and growth. This is easy to determine, because both processors and memory have been following their own empirical geometric scaling laws (the so-called Moore's laws) for the past twenty years, albeit with different geometric slope exponents—eighteen months to double for processor clocks and roughly twice that for the size of commodity DRAM—see Figure 1. The Semiconductor Industry Association (SIA) has projected continued scaling for another ten years, at least at the chip level, as shown in Figure 2 [6]. Recent fundamental work on the ultimate failure of MOSFETS by quantum tunneling through ultra-thin gate oxide layers supports this claim [7]. This scaling will total thirty years, a remarkable feat. This implies that faster CPUs, wide and fast memory buses, and fast system buses will continue to challenge the board designer for another decade. How much of a problem is this?




Figure 1. Actual scaling of the relative performance of CPU and main memory for commodity Silicon VLSI during the past two decades.




Figure 2. Projected scaling of clock frequency for commodity and high-performance CPUs based upon projections for Silicon VLSI by the Semiconductor Industry Association.

Before examining board-level optical interconnect opportunities, we first want to dispose of an application of great interest to the academic community but which is probably not practical and, in any case, is not required, except in specialized, low-volume (custom) environments. This is the on-chip (intra-chip) optical interconnect. SIA scaling shows that semiconductor processor and memory manufacturers expect to scale their technologies in density and speed for another ten years. In our view, this entirely eliminates on-chip optical interconnects from consideration for at least ten years and probably forever. If the reader doesn't accept this argument, then consider that adding lasers, which are large (1-10 micron) devices generating heat and noise, to VLSI circuits (which are sub-micron, densely packed circuits) might only make a hard thermal management and noise problem worse. Engineers are already dealing with devices dissipating nearly 100 W, as shown in Figure 3. Most of the power is dissipated in logic, not line drivers. Adding lasers makes this worse, not better.

The argument can be made, correctly [8,9], that scaling has entered a regime where the on-chip wiring is scaling unfavorably as feature size is reduced. We do not believe that this will cause the introduction of on-chip optical communications. The early impact will be mitigated by copper

wiring, additional metal layers, and low-K interlayer dielectrics. As time progresses we then expect to see more aggressive use of repeaters, architectural changes—multiple copies of registers, local control flow computation, and repartitioning of the system—that reduce the fraction of long wires, asynchronous circuit design styles that eliminate the need for low skew clock distribution, and extreme effort from the Silicon processing community to introduce thin film transistor repeaters within the metal system. For the most extreme requirements, cryogenic cooling gains a factor of seven reduction in the resistance of pure metal wiring without resorting to High Temperature Superconductors.




Figure 3. Power, voltage, and current curves over time for high-density, state-of-the-art logic chips. The modern CPU looks electrically more like a light bulb (power), an oven(current), or a flashlight (voltage), and will likely require advanced thermal management solutions.

Although some studies have indeed shown that optics can enhance VLSI performance, particularly in uncommon architectures [10,11,12,13,14], the integration of optics onto Silicon VLSI is fraught with fatal practical problems. The first problem is that no supplier of several hundred million transistor VLSI is going to ship a bare tested die to a supplier who will integrate optics onto the die and then retest, package, and sell the device. None of the present major Silicon

suppliers have manufacturing capability, or plans for capability, to integrate optics onto their die. Indeed, all of the mainstream suppliers instead believe that the SIA projections are targets they must achieve: the industry will squeeze every last transistor out of Silicon until gate oxides are so thin that quantum tunneling ruins the devices. Even then, the development of new (presently unknown) transistor designs are more probable than falling back on optics. Beyond the known good die problem lies the thermal management problem, the reliability problem, and finally the ‘this technology is too weird and I don’t understand it’ problem. These are formidable barriers to optoelectronic integration with VLSI.

First, CPUs are already nearing air cooling limits, and—despite limiting case arguments for lower power dissipation for optical signaling—optics for on-chip signaling will probably add net power to the chip. However, this problem could vanish if optical devices and their drivers are matched to the VLSI itself so that the entire link—VLSI-to-VCSEL-to light pipe-to receiver-to-VLSI—is optimized. This requires advances in design, manufacture, and test, all of which are possible but require fundamental changes in the VLSI industry.

Next, adding lasers, even short-wavelength VCSELs to VLSI, effectively doubles the failure rate of the integrated circuit. A typical CPU has a reliability of under 1000 FITs. A 100 element VCSEL array with a low 10 FITs per VCSEL has comparable reliability. Moreover, VCSEL lifetime is highly temperature activated, and it is difficult to see how an array will prove reliable when it is next to a 70C CPU die with local hot spots of order 100 C or more. Again, this problem is mitigated by optimizing the design of the optoelectronic system which includes the CPU, the optics, and the packaging.

Third, Silicon systems designers have CAD tools for transistors and gates, but not for lasers and photodetectors. Though advances in optical CAD are significant, their merge with commercial, large-scale chip and board layout CAD tools requires major investment, dedication, and time.

At present, bare VLSI devices provide a poor environment for VCSEL and receiver die and the two technologies—Si VLSI and GaAs lasers—don’t match up well in size, performance, and latency. As VLSI scales to 70 nm features within the decade, we’re led to ask, why break a technology with optics when it’s already predicted to continuing working with wires?

The further development of Silicon-based optoelectronics might seem to pose new opportunities for full optoelectronic integration. Indeed, it has been a dream of the optical interconnection field that the entire system—circuitry, sources, waveguides, and receivers—could be fabricated in a single material system. If this system is not GaAs—which lost out to Si for mainstream computer applications—could it be Silicon itself? Light sources are clearly the issue, as Silicon lasers are not on the horizon due to the indirect Si bandgap. But optical engineers are quite clever, and perhaps we may look forward to other innovative light sources or other methods to achieve (mostly) Si integration.

We return now to the interior of a computer box and look for optical interconnect opportunities. Most modern multiprocessor systems are cabinets containing multiple system boards—usually connected in a bus or switch configuration on a centerplane or backplane—which themselves contain multiple processors (perhaps 1-8), large quantities of fast SRAM and slower DRAM, and other integrated circuits, including communications ASICs, memory controllers, video chips, and glue logic [15,16]. Processor-to-memory interfaces are fast (Gbps), wide (100's of lines), and short (few cm), and are mainly fabricated as printed circuit board (PCB) traces or, in special cases, on multichip substrates.

Latency to main memory is a primary limitation on system performance. This latency is composed of a number of components including wire delay, address translation delays, arbitration delays, and the raw latency of the DRAMs themselves. In commercial, large scale systems, wire delay is a small component of the total. Adding optics to this link does not solve any technical problems, but it does add power dissipation and complexity to an already stressed and dense environment, and it increases latency by adding clock cycles (for the time required to get through the active driver and receiver chips). Moreover, the cost requirements, about \$.05 per line, eliminate optical technologies as a practical solution for on-board application [17]. Again, our remarks apply to commercial, mainstream systems such as servers, mainframes, and even high-performance graphics workstations. One can probably always create a custom environment for which optical interconnects at the on-chip and chip-to-chip level appear practical. The 'custom' usually implies 'at any cost' and is often used for unique, one-of-a-kind applications. But for mainstream VLSI systems, optics for intra-chip and inter-chip interconnects are not likely to be implemented for another ten years, if ever.

Continuing the analysis of in-the-box interconnects, we move up one level from on-board interconnect to board-board interconnects. Here we find our first commercial opportunity for optics within the box, and in fact, we find a very ripe opportunity. Systems boards usually run at some fraction of the processor clock, perhaps half as fast. Off-board communication speeds will thus approach 1 GHz in a few years. This poses the first real challenge for low-cost PCB design, namely, the routing of hundreds of 1 Gbps signals over 10-40 cm at 2.5 V or 1.8 V swings. Routing signals with these fast edge rates places extreme requirements on the electrical quality of the lines. Crosstalk and reflections must be adequately controlled for every signal.

Board designers might solve their problems by making use of large numbers of layers in their PCB to open up the average line-to-line spacing, or by resorting to differential signaling to improve noise rejection at the receiver and reduce ground current noise generation at the transmitter. The former solution dramatically increases cost and, after a few more years of clock doubling, ultimately fails. The latter solution doubles the number of pins on a communications ASIC and may result in intractable routing issues, particularly in the escape region under each high pinout package.

ASICs, with thousands of pins, are presently used as packet engines in shared memory systems; doubling the pin count is not a viable long-term option. Thus, the scaling predicted for Silicon will eventually drive the board designer to investigate new types of packaging and system design. Multichip packages aren't the answer for the majority of the market. But optical interconnections using free-space area arrays may be a viable answer. It is this solution we investigate here. We note that recent attempts to connect boxes with free-space links have proven practical [18].

Consider first the advantages of a free-space area interconnect to our hypothetical, stressed board designer. We'll deal with the realities of manufacturing such a link shortly. To be specific, consider an array interconnect of order 16x16 bi-directional lines, and assume each source/receiver is pitched at 1 mm on our interconnect die. This module has, in principle, 256 inputs and outputs that may be placed anywhere on the board by the designer, and it's under 4 cm². Locating these devices close to the major traffic points on the system boards can relieve the need to route large numbers of high speed signals to the traditional edge connector. For applications requiring only board-to-board interconnection, an area array of this type may be

located anywhere on the boards and, in particular, next to the transmitting and receiving ASICs. This eliminates the problem of running high-speed, single-ended or differential lines, from the ASIC to the board edge. It replaces it with the problem of designing and building these optical links and obtaining and maintaining alignment between two parts that have a few cm separation but require in principle sub-mm alignment. It should be clear here that scaling this interconnect to several Gbps per line is, in principle, the same problem as is being solved by VCSEL link suppliers for multi-Gbit serial links or parallel fiber-optic links. Since CPUs are projected to run at many GHz in the next decade, but not faster, board ASICS will run at comparable rates, and area VCSEL links can provide the interconnection bandwidth required in the coming decade.

Links of this type can also be used as very high performance replacement for current large edge connectors. This can improve system noise margin because crosstalk and ground noise coupling become more difficult to control in traditional connectors as edge rates increase. An additional advantage is galvanic isolation that leads to a natural 'hot-plug' capability.

There remain two problems. The first is designing and building an array of VCSELs, receivers, or hybrids of lasers and photodetectors on single assemblies that can modulate at many GHz with suitable electrical isolation and minimal crosstalk. Some success in optical monolithic integration (lasers and lenses, lasers and receivers) has already been achieved [19,20,21], and a number of efforts have successfully integrated CMOS with VCSEL devices [22,23]. It should be clear that this is not a fundamental limitation, since multi-GHz Silicon VLSI is tractable. Most fundamental issues for arrays have been studied or demonstrated. Hence, what remains is to perform an exact system design and build. Sophisticated CAD tools for free-space optical systems are already in development [24].

A potentially more difficult problem is maintaining the alignment of free-space links [25]. A practical system must function when the manufacturer or repairman inserts boards that snap into guides with a few mm tolerance, and must stay aligned as the system heats up to its equilibrium operating point. To achieve lateral alignment, the individual beams may be expanded by using either external or integrated lenslet arrays. However, it is well known that nothing is free, and that the engineer is trading off lateral alignment tolerance for angular alignment tolerance when using expanded beams. Alignment can be achieved by using active control systems based on

micromechanical technologies, or by ingenious invention. However it is achieved, laser alignment is a practical, not a fundamental issue, that can be solved by creative engineering. Multiple boards will clearly require some repeater modules on each individual board, but these may be constructed in the same fashion as the links.

Each computer, telco switch, router, and network switch designer will have his own unique board, center- or backplane, and cabinet design; a supplier of free-space optical interconnect (FSOI) modules will face a nightmare attempting to supply custom parts to each unless a standard form factor is selected. We believe that the computer systems industry and the component supply industry ought to attempt to standardize an FSOI form factor. When this is achieved, free-space optics will have broken the barrier separating the physical layer of the network from that of the computer, and scaling should continue well into the next decade.

IV. Conclusions

With this analysis, we now have a complete hierarchy of interconnects for computing machines for the next ten years. Optical interconnects are already external to the box and are here to stay. Ironically, free-space optical interconnects - which have yet to be commercialized - are most promising for use inside the system to connect boards in an areal fashion. Essentially all of the technology has been demonstrated to achieve this. What remains is a concerted effort by industry to develop products and a standard whereby these products will be used.

V. References

- [1] Michael J. Yadlowsky, Evelyn M. Deliso, and Valeria L. Da Silva, "Optical fibers and amplifiers for WDM systems," Proc. IEEE 85, 1765 (1997).
- [2] Information on the various interconnect standards mentioned in the text may be found at the following locations: Fibre Channel Association, <http://www.fibrechannel.com/>;
Gigabit Ethernet, <http://www.gigabit-ethernet.org/>; High-speed study group, http://grouper.ieee.org/groups/802/3/10G_study/public/june99/index.html;
- [3] Kent D. Choquette and Hong Q. Hou, "Vertical-cavity surface emitting lasers: Moving from research to manufacturing," Proc. IEEE 85, 1730 (1997).
- [4] Product literature is available at www.infineon.com
- [5] Product literature is available at www.optobahn.com
- [6] "The National Technology Roadmap for Semiconductors: Technology Needs," Semiconductor Industry Association (1997).
- [7] D.A Muller, T. Sorsch, S. Moccio, F.H. Baumann, K. Evans-Lutterodt, and G. Timp, "The electronic structure at the atomic scale of ultrathin gate oxides," Nature 399, 758 (1999).
- [8] D.A.B. Miller, "Physical reasons for optical interconnection," Special Issue on Smart Pixels, J. Optoelectronics 11, 155 (1997).
- [9] Evan E. Davidson, Bradley D. McCredie, and Walter V. Vilkelis, "Long lossy lines and their impact upon large chip performance," IEEE Trans. Comp., Pkg., and Mfg. Tech. 20, 361 (1997).
- [10] Gokce I. Yayla, Philippe J. Marchand, and Sadik C. Esener, "Speed and energy analysis of digital interconnections: Comparison of on-chip, off-chip, and free-space technologies," App. Opt. 37, 205 (1998).

- [11] Ashok V. Krishnamoorthy and Keith W. Goosen, "Optoelectronic-VLSI: photonics integrated with VLSI circuits," *IEEE Spec. Top. Quant. Elec.* 4, 899 (1998).
- [12] Ashok V. Krishnamoorthy and David A.B. Miller, "Scaling optoelectronic-VLSI into the 21st century: A technology roadmap," *IEEE J. Sel. Top. Quant. Elec.* 2, 44 (1996).
- [13] Osman Kibar, Daniel A. Can Blerkom, Chi Fan, and Sadik C. Esener, "Power minimization and technology comparisons for digital free-space optoelectronic interconnections," *IEEE J. Lightwave Tech.* 17, 546 (1999).
- [14] Ashok V. Krishnamoorthy and David A.B. Miller, "Firehose architectures for free-space optically interconnected VLSI circuits," *J. Parallel and Distributed Computing* 41, 109 (1997).
- [15] Alan Charlesworth, "Starfire: Extending the SMP envelope," *IEEE Micro*, pp. 39-49 (Jan/Feb 1998).
- [16] Alan Charlesworth, Andrew Phelps, Ricki Williams, and Gary Gilbert, "Gigaplane-XB: Extending the Ultra server family," *Proc. Hot Interconnects V*, pp. 1-16 (Sep 1997).
- [17] Ashley Saulsbury, "Why computer architects don't need to see the light," *IEEE Workshop on interconnections within high-speed digital systems*, Santa Fe, NM (May 1998).
- [18] Dominic J. Goodwill, "Free-space optical interconnect for terabit network elements," *Proc. Optics in computing conference*, Snowmass CO (April 1999).
- [19] D.A. Louderback, "Modulation and free-space link characteristics of monolithically integrated vertical-cavity lasers and photodetectors with microlenses," *IEEE J. Sel. Top. Quant. Elec.* 5, 157 (1999).
- [20] O. Sjolund, D.A. Louderback, E.R. Hegblom, J. Ko, and L.A. Coldren, "Monolithic integration of substrate input/output resonant photodetectors and vertical cavity lasers," *IEEE J. Quant. Elec.* 35, 1015 (1999).

- [21] Peter Kazlas, John Wasserbauer, Jeffrey Scott, David Paananen, Stan Swirhun, and David Lewis, "Monolithic vertical-cavity laser/p-i-n photodiode transceiver array for optical interconnects," *IEEE Phot. Tech. Lett.* 10, 1530 (1998).
- [22] A.V. Krishnamoorthy, W.S. Hobson, R.E. Leibenguth, S.P. Hui, G. J. Zyzdik, K.W. Goosen, J.D. Wynn, B.J. Tseng, J. Lopata, J.A. Walker, J.E. Cunningham, and L.A. D'Asaro, "Vertical-cavity surface-emitting lasers flip-chip bonded to gigabit-per-second CMOS circuits," *IEEE Phot. Tech. Lett.* 11, 128 (1999).
- [23] T.K. Woodward and Ashok V. Krishnamoorthy, "1-Gb/s integrated optical detectors and receivers in commercial CMOS technologies," *IEEE J. Sel. Top Quant. Elec.* 5, 146 (1999).
- [24] Steven P. Levitan, Timothy P. Kurzweg, Philippe J. Marchand, Mark A. Rempel, Donald M. Chiarulli, Jose A. Martinex, John M. Bridgen, Chi Fan, and Frederick B. McCormick, "Chatoyant: a computer-aided design tool for free-space optoelectronic systems," *App. Opt.* 37, 6078 (1998).
- [25] David T. Neilson, "Tolerance of optical interconnections to misalignment," *Appl. Opt.* 38, 2282 (1999).

About the Authors

Rick Lytel is a Distinguished Engineer at Sun Microsystems, where he directs the physics department of Sun Microsystems Laboratories. His research interests include the application of optical interconnection to high-performance computer systems, potential applications of molecular electronics, applications of small world networks to interconnection, and thermal management technologies.

Rick was a founder of ROI Technologies, now Lightwave Microsystems, and also founded and was VP and General Manager at Akzo Nobel Photonics. He also started and ran the Lockheed Photonics Technology group, where he and his group started the field of nonlinear optical polymer devices. He is an adjunct Professor of physics at Washington State University and lectures annually on optical interconnects, computer systems, and other subjects. He has published over 70 papers.

Howard L. Davidson is currently a Distinguished Engineer at Sun Microsystems laboratories with responsibility for advanced physical technology. Current research areas include high power density cooling, optical interconnects, molecular electronics, quantum computing and SETI.

He has previously held product development and research positions at Schlumberger Palo Alto Research, Lawrence Livermore National Laboratory, CRAY Research, and Hewlett-Packard. Previous research areas include integrated circuit testing, supercomputer design, ultra-light satellites, high thermal conductivity, thermal materials, and electronic instrumentation.

Dr. Davidson is a recognized authority on physical system design. He has chaired the IEEE Computer and System Packaging Workshop and the IEEE Workshop on High Speed Interconnections within Digital Systems. He has served as an Associate Editor for IEEE Transactions on VLSI Systems. Dr. Davidson has presented many papers, invited papers and tutorials on high speed digital system design, thermal management, and optical interconnects for digital systems. He is the author of the chapter on inspection and testing in "Multichip Module Design, Fabrication, & Testing". Dr. Davidson holds 16 patents. He is a Senior member of IEEE, and a member of APS and AAAS. Howard L. Davidson has Ph.D. in Physics from Wesleyan University.

Nyles Nettleton is a staff engineer at Sun Microsystems Laboratories in Mountain View, California. During his career at Sun he has worked on medical imaging, microprocessor design, graphics accelerators and algorithms, systems performance analysis, clustering technologies, disk arrays and supercomputer design. His research interests include systems architecture and design, high-speed communications, optics and nanotechnology.

Before joining Sun in January of 1988, Nyles worked at Arete on MC68020 based multiprocessor transaction processing platforms. Before that, he worked at Silicon Graphics, where he worked on the Iris 1000 graphics workstation.

Theresa Sze is a member of technical staff at Sun Microsystems Laboratories in Mountain View, California, where she is designing optical connectors for servers. Her research interests include optical systems, high-speed systems, server architecture, and mobile computing systems.

Before joining Sun, Theresa worked on research projects at University of Illinois, Fermi National Laboratory and University of New Mexico, from which she received her M.S. in Electrical Engineering.