

Optically Interconnected Parallel Computing Systems

Researchers have developed an architecture for high-speed computation, image processing, and robotic vision systems that uses both the programmability of mature electronic technology and the density and parallelism of high-speed optical interconnects.

**Masatoshi
Ishikawa**

Neil McArdle
University
of Tokyo

As silicon electronic processors achieve greater integration density and faster clock speeds, data communication requirements become more demanding. Electrical bandwidth is inherently limited by capacitive and inductive effects at high frequency, electromagnetic interference, and pin-density requirements at the chip and board level. All these things seriously inhibit conventional interconnection technologies.

Tomorrow's systems will need high-bandwidth and dense communication paths at various levels:

- Intercabinet connections among individual computer systems.
- Backplane-level connections among boards.
- Board-level connections among multichip modules (MCMs) on boards.
- Chip-to-chip connections on a board or MCM.
- Gate-to-gate connections within a chip.

Commercial high-performance computers are now beginning to use optical interconnections at the inter-cabinet level. These connections usually consist of optical fiber ribbons, with each fiber carrying signals at 1 to 2 Gbits per second over distances of 200 to 300 meters. The aggregate bandwidth is as much as 30 Gbits per second.

Some manufacturers of specialized high-performance systems are also investigating optical technologies at lower levels of the hierarchy. For example, Cray Research is studying the use of polymer optical wave guides to distribute the optical clock in its T-90 computer.¹ Mercury Computer Systems is also investigating optical backplane technologies, which it believes will be essential for some of its high-performance systems.²

Compared to VLSI fabrication, interconnection technologies have advanced slowly. Conventional tech-

nologies will have difficulty meeting the future requirements for communication among processors and between a processor and memory. The Semiconductor Industry Association predicts that in the year 2007, systems will require as many as 5,000 I/Os, with off-chip clock rates of about 500 MHz.³

However, if we integrate suitable optoelectronic devices with silicon electronics, we can use optical communication channels to transfer data on and off chips. Optics can effectively communicate data to the chip surface in a massively parallel fashion and at high speed. Optical connections are not limited to wire bonds around the edges of the chip die. Nor do they suffer from electromagnetic interference and high frequency losses. In addition, they can cross in free space without experiencing signal crosstalk among channels. And once the signal is in the optical domain, the driving power is independent of interconnection distance, given no attenuation losses.

At the University of Tokyo, we have been working on some technologies to overcome the interconnection bottlenecks experienced by high-speed parallel processing systems. In particular, we have developed an optically interconnected architecture for high-speed computation, image processing, and robotic vision systems. Our architecture utilizes the programmability of mature electronic technology while taking advantage of the density and parallelism of high-speed optical interconnects.

OPTICALLY INTERCONNECTED SYSTEMS

There are two main approaches for realizing optical connection paths:

- *Active optical devices are integrated with silicon electronics and are used to emit the signal beams.* These devices include light-emitting diodes (LEDs) and vertical cavity surface-emitting lasers.

The sidebar, "Other Optical Computing Research in Japan," describes a project in VCSELs.

- *Modulating devices that vary the intensity, phase, or polarization of an incident optical beam, according to an electronic signal.* This category includes symmetric self-electro-optic-effect devices (S-SEEDs). Light from an external source is directed to these devices, and the S-SEEDs modulate and reflect the beam to another optoelectronic chip. This approach requires some extra optical components and a laser source.

Both types of devices are typically fabricated using multiple layers of compounds based on gallium arsenide (GaAs). They then are combined with the silicon electronics to make a *smart-pixel device*.

One way to combine them is called *solder-bump flip-chip bonding*, in which the silicon and GaAs devices are joined by an array of small solder bumps (each about 100 μm in diameter) distributed across the surface of one of them. Some success has been achieved in this area for S-SEED modulators on CMOS (complementary metal-oxide semiconductor) electronics.⁴ Solder-bump flip-chip bonding may also eventually be applied to VCSELs and silicon electronics, although demonstrations using these devices are not completely developed.

Another way to combine optical I/O devices and processing circuitry is to monolithically integrate all devices on a single GaAs chip. This approach is briefly described in the sidebar, "Other Optical Computing Research in Japan."

Demonstration systems

Most optoelectronic demonstration systems target special-purpose applications and use S-SEED modulators on silicon CMOS technology. The most popular application targets are fast Fourier transforms (FFTs) and sorting. Applications like these have inherently global interconnection requirements, in which data nodes or processing elements must communicate with many distant nodes.

An FFT machine currently under investigation at the University of North Carolina, Charlotte, uses custom silicon CMOS circuitry and GaAs modulators. It exhibits an I/O bandwidth of 29 Gbytes per second and can calculate a 1,024-point FFT in a few microseconds.⁵ Sorting machines are under construction at Colorado State University and Heriot-Watt University in Scotland.⁶ The latter machine uses two optically interconnected smart-pixel arrays interconnected through a *perfect-shuffle* optical system, which interleaves two halves of the array like two halves of a deck of cards. It achieves a full sort of 1,024 16-bit words in about 10 μs .

Other Optical Computing Research in Japan

Much of the optical computing research in Japan is performed under the auspices of the Real World Computing-Real World Intelligence/Parallel and Distributed Computing (RWC-RWI/PDC) project, a program sponsored by the national government that involves industry and university partners.

This research focuses on next-generation computing systems, including parallel and distributed computing, algorithmic theory, optical interconnections, and artificial intelligence. Work in optical interconnections is being pursued by Matsushita, NEC, Hitachi, Oki, Mitsubishi, and the University of Tokyo.

Nippon Telegraph and Telephone is active in the monolithic integration of processing circuitry and optical sources.¹ In this, the company combines a photodetector, a vertical cavity surface-emitting

laser (VCSEL), and a field-effect transistor in a single pixel, all fabricated in gallium arsenide (GaAs). Thus, no silicon-GaAs bonding techniques are required. However, the processing circuitry is presently limited to a single switching transistor and does not utilize the breadth of silicon CMOS electronics.

A crucial device for optoelectronic computing systems is the VCSEL. Among the leading research groups in this field is NEC. VCSELs have numerous advantages for optical interconnection systems. Since they emit light normal to the chip surface, many lasers can be integrated into a single chip to achieve dense two-dimensional arrays—for example, up to 8×8 .² They have relatively low threshold currents (of the order of 1 mA), and recent advances have reduced this to the order of tens of microamperes.³ The output light can be in a single spatial mode (Gaussian beam) and thus can be efficiently coupled into optical

fibers and small photodetector windows.

Current VCSELs also have relatively small array sizes (a maximum of 8×8), which limits their use in some applications.

References

1. S. Matsuo et al., "Monolithically Integrated Photonic Switching Device Using an MSM PD, MESFETs, and a VCSEL," *IEEE Photonics Technology Letters*, Vol. 7, 1995, pp. 1,165-1,167.
2. T. Yoshikawa et al., "Complete Polarization Control of 8×8 Vertical-Cavity Surface-Emitting Laser Matrix Arrays," *Applied Physics Letters*, Vol. 66, 1995, pp. 908-910.
3. G.M. Yang, M.H. MacDougal, and P.D. Dapkus, "Ultralow Threshold Current Vertical-Cavity Surface-Emitting Lasers Obtained with Selective Oxidation," *Electronics Letters*, Vol. 31, 1995, p. 886.

Architectural consequences

The use of *surface-normal* optical interconnections—that is, optical interconnects distributed on the surface of a chip—has some interesting consequences for system architectures. The highly parallel two-dimensional I/O that is achieved can be applied to high-speed page-oriented optoelectronic memories or to an interface to volume optical memories, such as holographic memories.⁷ A page of data can be read or written in a single clock cycle, and the data rate is not limited by the bus bottleneck between the processor and memory.

Surface-normal optical interconnections are also ideally suited for pipelined optoelectronic systems that process images or two-dimensional patterns. The University of Southern California is building a parallel-pipelined, smart-pixel, image-processing system that is particularly suited to fast image-compression techniques.⁸

One area of research in our laboratory in the Department of Mathematical Engineering and Information Physics at the University of Tokyo is real-time image processing for machine vision. Vision systems that use conventional cameras encounter a serious bottleneck when they try to convert two-dimensional images into a serial signal stream for transfer to the processing unit. The frame rate is usually limited to standard video frame rates of around 30 frames per second, especially when there are large numbers of pixels. In our work we have sought to remove the bottleneck by integrating processing and optical I/O elements. This achieves high-speed image processing with frame rates of a few microseconds, which is especially useful for high-speed target tracking and robotic vision.

We have found that a local processing element (PE) placed at each sensor in the camera (or photodetector [PD] array) removes the need for parallel-to-serial conversion and scanning techniques. We use an integrated array of programmable PEs and PDs as a smart sensing element. By stacking many of these arrays in a pipelined system—with a free-space optical interconnection between them using a VCSEL array—we can utilize the global connectivity of the optics and the processing functionality of electronics to build a general-purpose machine for various applications. Similar work at the University of Colorado at Boulder uses silicon electronics and VCSELs for a general-purpose parallel optoelectronic processor.⁹

SYSTEM ARCHITECTURE AND DESIGN

Our architecture uses the pipelined configuration shown in Figure 1.¹⁰ Each chip contains a two-dimensional array of PEs, and each PE contains some processing circuitry, registers, and local memory. PEs communicate with their four nearest neighbors via on-chip electrical connections. Integrated with each PE is




Figure 1. Pipelined, optically interconnected architecture.

a PD for optical input and an optical output device, such as a modulator, LED, or surface-emitting laser diode.

This architecture differs from conventional parallel processing systems in three ways:

- It supplies inputs and outputs in parallel.
- It supplies optical paths to connect nonlocal processors.
- Its interconnection topology is reconfigurable: The light beams can be dynamically redirected.

The first two features supply high connection density and global connectivity. Although low-level image-processing tasks generally use neighborhood connections, some high-level tasks, such as moment extraction and feature extraction, can benefit from global interconnection paths. Global optically interconnected parallel-processing systems are also suitable for computation-intensive applications, including sorting, FFTs, signal processing, matrix operations, and high-level image processing, where the nature of dataflow is inherently nonlocal.

INTEGRATED OPTOELECTRONIC PROCESSING DEVICE

Our first-generation optoelectronic processing device, the SPE-8 (Sensory Processing Elements 8), contained a linear array of eight PEs. Each PE had optical I/O connections to a phototransistor and an LED, as well as connections with neighboring PEs. Many SPE-8 chips, therefore, could be connected to form a large two-dimensional array of parallel processors.

The SPE-8 allowed us to construct several demonstration systems, summarized in Table 1. The SPE-4k had no optical interconnections: It was an experi-

Table 1. Optoelectronic processing devices and systems.

	Devices		Systems		
Feature	SPE-8 (1992)	S3PE (1995–1997)	SPE-4k (1993)	SPE-II (1994)	Ocular (1997)
Architecture	General-purpose processing	General-purpose processing	Neighbor-connected image processor	Optically interconnected feedback system	Optically interconnected feedback system
PE	1 × 8 digital; 1,348 transistors/PE	8 × 8 digital; 437 transistors/PE	64 × 64 digital; using SPE-8	8 × 8 digital; using SPE-8	8 × 8 digital; using S3PE
Interconnection Optics			1 to 1 optical input/output	Reconfigurable 1 to n	Reconfigurable 1 to n
Electronics	4 neighbors	4 neighbors	4 neighbors	4 neighbors	4 neighbors
Optical source			Discrete LED 8 × 8	VCSEL 8 × 8	VCSEL 8 × 8
Optical input			Discrete photo- transistor array	8 × 8 Si PD array	PE and PD integrated chip (S3PE)
SLM				PAL-SLM (CGH)	Compact PAL-SLM (CGH)
Realization form	8 PEs per LSI gate array	Full custom CMOS VLSI	Scaled-up system (image processing)	Discrete units prototype (parallel processing)	Compact integrated system (parallel processing)

mental image-processing system that contained 64×64 PEs constructed from 512 SPE-8 chips. Each PE was connected directly to a phototransistor, and the output from the PEs drove an LED array.

The SPE-4k demonstrated various real-time, low-level image-processing algorithms: edge detection, skeletonization, moving-object detection, and trace. Since there were no parallel-to-serial conversion bottlenecks, the SPE-4k could process images about 10^5 times faster than conventional image-processing systems, which are limited to about 30 frames per second.

Higher computing performance can be achieved by introducing dense parallel optical interconnections between the PE arrays, as Figure 1 shows. These optical paths can provide global connectivity, which would

require many iterations using the neighborhood electrical connections. One system, SPE-II, demonstrates this architecture. SPE-II uses a single layer of PEs in a feedback-type configuration, in which the optical outputs from the PE array are directed back onto the PDs of the same PE array, thereby allowing data to circulate. The system contains an 8×8 array of PEs constructed from eight SPE-8 chips.¹⁰ The outputs from the PEs drive an array of 8×8 surface-emitting lasers, from which the light is imaged to a PD array connected to the inputs of the PEs. The data, represented by intensity-modulated laser beams, flows around the system in a cyclical fashion. The operation of the PEs change with each cycle, according to the desired algorithm.

SPE-II's optical interconnection paths come with some degree of programmability. By inserting a computer-generated hologram (CGH) into the optical paths between the lasers and the PDs, the system can deflect the beams arbitrarily, thus altering its interconnection topology. In SPE-II, we implemented a dynamic CGH using a liquid-crystal spatial light modulator (SLM), a device for introducing a two-dimensional pattern on an incident laser beam by modulating its intensity, phase, or polarization.¹¹ It can change the CGH dynamically under the control of a host computer. Thus, the SPE-II was a highly general-purpose processing machine.

In later generations, we drastically reduced the size of these systems by integrating the PD at each PE into a single VLSI chip. Not only did this drastically reduce system size, it removed the difficulty of wiring large, dense arrays of PEs and PDs in parallel. There are some trade-offs involved in the design of the PEs: High-functionality PEs need more transistors and are therefore larger, but small PEs are needed to make high-resolution arrays on a single chip. Our architec-

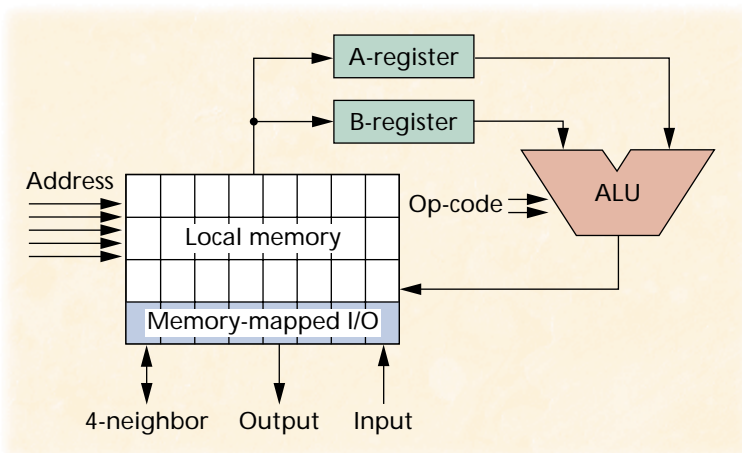


Figure 2. Internal architecture of one processing element of S³PE, a general-purpose bit-serial processor with local memory.

ture calls for a general-purpose PE; however, for vision systems a large number of pixels is required, so the pixel size must be kept small.

The S³PE (smart-and-simple sensory processing element), shown in Figure 2, is an integrated single-chip 8 × 8 array of PEs and PDs. The S³PE contains an arithmetic logic unit (ALU) that can perform bit-serial operations, and it has connections to neighboring PEs and I/O for the optoelectronic interface. Figure 3 shows a layout of the S³PE. It is relatively compact (around 390 μm × 290 μm), yet it is programmable and has some local memory. We use this device in our latest demonstrator system, Ocular.

OCULAR

Ocular (Optoelectronic Computer Using Laser Arrays with Reconfiguration) uses the S³PE device coupled to a VCSEL array, along with compact optics and optomechanics. As in SPE-II, it uses a single S³PE array in the feedback configuration, shown in Figure 4. The outputs from the VCSEL array are imaged through the optical system onto the PDs. Between two gradient




Figure 3. Layout of S³PE chip with 8 × 8 PEs and PDs.

index (GRIN) lenses is a reflecting SLM on which is displayed CGHs. The operation executed by the PE array can be changed on each cycle, and the interconnection topology can also be changed by updating the CGH.





Figure 4. Experimental implementation of the Ocular system.

Figure 5. Ocular optomechanical structure (dark base) with GRIN rod lenses (light cylinders).



The PDs of the present S³PE chip are approximately $20 \times 20 \mu\text{m}^2$. The chip also has some simple receiver and amplifier circuitry, but the receiver speeds are relatively slow (a few MHz) in these first prototypes. An important issue in smart-pixel devices—where electronic processing, detectors, and high-speed receivers must be compactly integrated—is power dissipation. Many researchers are investigating new receiver designs that are suitable for smart-pixel systems.¹²

Connecting two-dimensional arrays of optoelectronic devices requires high-performance and cost-effective imaging systems since the imaging requirements are crucial to overall system performance. Typical requirements in smart-pixel systems include being able to image devices of the order of $5\text{--}20 \mu\text{m}$ and having device pitches of $20\text{--}250 \mu\text{m}$ in arrays of 8×8 to 64×64 . In addition, the lenses must be compact and inexpensive.

When we compare the resolution, information density, compactness, complexity, and cost of two types of imaging systems of similar diameter—one using conventional off-the-shelf achromatic lenses, the other using GRIN rod lenses—we find that GRIN lenses cost less and have up to five times greater capacity in useful number of channels that can be imaged.¹³ GRIN rods have been successfully used in parallel interconnection systems in our laboratory and elsewhere.¹⁴

Optomechanics and packaging of the system are also important issues. We have designed a custom optomechanical system, shown in Figure 5, that satisfies the requirements of compactness, mechanical and thermal stability, reliability, ease of fabrication and alignment, low cost, and modularity of design to enable easy extension of the system architecture. It consists of V-groove machined holders in which the

GRIN rod lenses are secured. The axes of the V-grooves are precisely machined and aligned prior to insertion of the lenses, thereby simplifying the alignment procedure after assembly. Using these in the Ocular system, we have achieved a reduction in the system volume of two orders of magnitude over the previous SPE-II system. Furthermore, extensive environmental and thermal cycling tests have shown the optomechanics to be extremely stable.

APPLICATIONS AND ALGORITHMS

Although many architectures have been proposed, very little research has been done on the implications of a dynamically reconfigurable interconnection topology. Architectures of this sort have many consequences for more efficient algorithm design and parallel programming.

Our laboratory is developing a theory by which an arbitrary application can be embedded into an optically connected hierarchical architecture of the type shown in Figure 1. We can analyze such systems from an algorithmic viewpoint and evaluate some basic tendencies, considering both the performance of optical interconnections and the application's calculation time. We believe that various functions can be performed by manipulating the optical interconnections among PE arrays. The functions include one-to-many interconnections (broadcast operations) and spatial-shift interconnections (summation operations). A typical parallel algorithm could be implemented by successively applying these operations. For example, matrix-vector multiplication, dynamic programming, and sorting algorithms can be implemented by combining broadcast and summation operations.

Our analysis of this theory has generated some interesting results. For example, we found that small optical interconnection fan-outs could improve the performance for typical algorithms,¹⁵ a result that affects the design of entire systems. For example, low optical fan-outs indicate that low optical output power—and consequently simpler optical receiver circuitry—are required. It also reduces the resolution requirements on the SLM since fewer pixels are required for a low fan-out CGH.

Consider a matrix-vector multiplication. To calculate the multiplication of an 8×8 matrix and an eight-element vector, the values of each of the elements of the matrix are preloaded into the PE array and stored in local memory. The vector is loaded into the base row of the PE array, and its elements have to be broadcast to each row of the matrix. We use the optical broadcast operation to perform this efficiently. A hologram that connects the eighth row to the fourth row is written on the SLM, as shown in Figure 6a. Next, a hologram is written that connects row 4 to rows 3, 2, and 1, and row 8 to rows 7, 6, and 5, as shown in

Figure 6b, such that all rows of the PE array contain the vector after two steps. The use of local electrical connections would have required seven steps to propagate the vector across the entire PE array. The multiplication of the array and the vector is then calculated by the ALU of each PE.

Next, using the neighborhood connections, each PE transmits its result to the PE on its right, where it is added to the data contained in that PE. This cycle is repeated three times, after which the PEs in column 4 contain the sum of columns 1 to 4, and the PEs in column 8 contain the sum of columns 5 to 8. Finally, the optical summation operation is applied by writing a hologram that connects column 4 to column 8, as shown in Figure 6c.

Although the frame rate of the SLM (70 Hz) is slow, this example demonstrates the efficient use of both local electrical connections and global optical connections, and the optical interconnection pattern is dynamically changed depending on the connections required by the algorithm. Some algorithms require optical reconfiguration at the clock rate of the PEs, others require reconfiguration occasionally (as in matrix-vector multiplication), while others require configuration only once before execution of a program. Further work needs to be done to determine the optimum reconfiguration rate for various applications and to find which devices have suitable frame rates.

We could reduce PE size by further optimizing its design and by using smaller CMOS fabrication processes. This would allow more complex PEs and the integration of more PEs in a single chip.

Some developments of VCSELs include lowering the emission wavelength (to better match silicon detectors), improving yield, decreasing the threshold current, and reducing heat dissipation so that larger arrays can be operated without thermally induced failure. It will be necessary in the future to integrate the S³PE chip with the VCSEL array. Some promising techniques, such as solder-bump flip-chip bonding, have been developed in this field and are nearing maturity.

By taking a modular approach in the design of the optics and optomechanics, we can easily extend our system to more complex architectures—for example, those consisting of several layers of PE arrays.

Some technical difficulties remain, but many of the required technologies are nearing maturity. Other technologies will require improvements in device performance so that densely integrated arrays of parallel processing circuitry can be combined with optical I/O for interprocessor communication. In order to take advantage of the ability of optical interconnects to reconfigure, more work needs to be done on the impli-

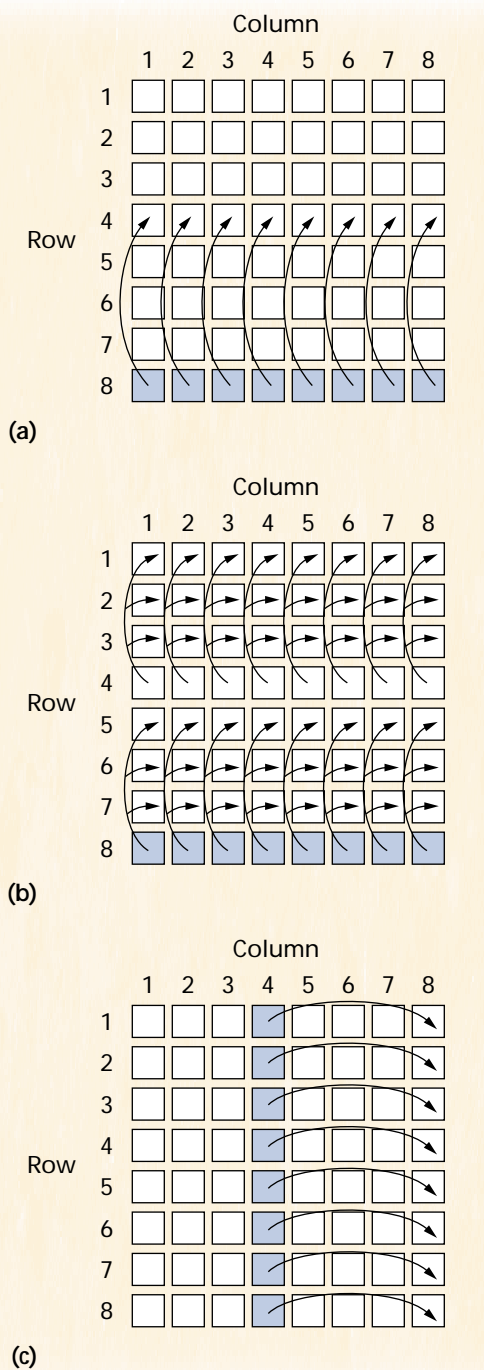


Figure 6. Matrix-vector multiplication using reconfigurable interconnections: (a) One CGH connects row 8 to row 4 optically. (b) For the multiplication to take place, a second CGH connects row 4 to rows 3, 2, and 1, and row 8 to rows 7, 6, and 5. (c) A third CGH connects column 4 to column 8 during the summation step.

cations and benefits of a dynamically programmable interconnection topology.

Optoelectronic parallel-processing systems will overcome some of the interconnection problems facing conventional electronic technology—allowing high-speed computers powerful enough for vision and image-processing applications. ♦

Acknowledgments

We thank Makoto Naruse and Takashi Komuro of the University of Tokyo. We also thank Kenichi Kasahara and his group at NEC Optoelectronics Research Laboratories for supplying the VCSEL arrays used in these experiments, and Tsutomu Hara, Yuji Kobayashi, and Haruyoshi Toyoda of Hamamatsu Photonics for supplying the PAL-SLM and the custom optomechanical system. We thank the Japan Society for the Promotion of Science for supporting Neil McArdle through a postdoctoral research fellowship.

References

1. R.T. Chen et al., "Si CMOS Process-Compatible Guided-Wave Multi-Gbit/sec Optical Clock Signal Distribution System for the Cray T-90 Supercomputer," *Proc. Massively Parallel Processing with Optical Interconnections '97*, IEEE CS Press, Los Alamitos, Calif., 1997, pp. 10–24.
2. C. Lund, "Optics Inside Future Computers," *Proc. Massively Parallel Processing with Optical Interconnections '97*, IEEE CS Press, Los Alamitos, Calif., 1997, pp. 156–159.
3. *The National Technology Roadmap for Semiconductors*, Semiconductor Industry Assoc., San Jose, Calif., 1994.
4. K.W. Goosen et al., "GaAs MQW Modulators Integrated with Silicon CMOS," *IEEE Photonics Technology Letters*, Vol. 7, 1995, pp. 360–362.
5. R.G. Rozier, F.E. Kiamilev, and A.V. Krishnamoorthy, "Design and Evaluation of a Photonic FFT Processor," *J. Parallel and Distributed Computing*, Vol. 41, 1997, pp. 131–136.
6. M.P.Y. Desmulliez et al., "Perfect-Shuffle Interconnected Bitonic Sorter: Optoelectronic Design," *Applied Optics*, Vol. 34, 1997, pp. 5,077–5,090.
7. M. Schaffer and P.A. Mitkas, "Smart Photodetector Arrays in Silicon CMOS for Page-Oriented Optical Memories," *IEEE Trans. Electron Devices*, to appear.
8. J.-M. Wu et al., "Smart Pixel Array Cellular Logic (SPARCL) Processor for Eliminating SIMD I/O Bottlenecks: System Demonstration and Performance Scaling," *Optics in Computing*, OSA Technical Digest Series, Vol. 8, 1997, pp. 152–154.
9. V. Morozov, J. Neff, and H.J. Zhou, "System Analysis of Global Free-Space Optical Interconnect," *Optics in Computing*, OSA Technical Digest Series, Vol. 8, 1997, pp. 224–226.
10. M. Ishikawa, "Optoelectronic Parallel Computing System with Reconfigurable Optical Interconnection," *Critical Reviews*, Vol. CR62, 1996, pp. 156–175.
11. A. Kirk, T. Tabata, and M. Ishikawa, "Design of an Optoelectronic Cellular Processing System with a Reconfigurable Holographic Interconnect," *Applied Optics*, Vol. 33, 1994, pp. 1,629–1,639.
12. T.K. Woodward, "VLSI-Compatible Smart-Pixel Interface Circuits and Technology," *Smart Pixels Technical Digest*, Aug. 1996, p. 65.
13. N. McArdle et al., "Experimental Realization of a Smart-Pixel Parallel Optoelectronic Computing System," *Proc. Massively Parallel Processing with Optical Interconnections '97*, IEEE CS Press, Los Alamitos, Calif., 1997, pp. 190–195.
14. A. Kirk et al., "Compact Optical Imaging System for Arrays of Optical Thyristors," *Applied Optics*, Vol. 36, 1997, pp. 3,070–3,078.
15. M. Naruse et al., "An Algorithmic Approach to Hierarchical Parallel Optical Processing Systems," *Optical Computing Technical Digest*, 1996, pp. 102–103.

Masatoshi Ishikawa is an associate professor in the Department of Mathematical Engineering and Information Physics, Faculty of Engineering, University of Tokyo. His research interests include optoelectronic computing, parallel processing, machine vision, sensor fusion, and robotics. He received a BS, an MS, and a PhD in engineering from the University of Tokyo. Ishikawa is a member of the Japanese Applied Physics Society.

Neil McArdle is a visiting research fellow in the Department of Mathematical Engineering and Information Physics, Faculty of Engineering, University of Tokyo. His research interests include optoelectronic computing, optical interconnections, and optical design. He received a BS in laser physics and optoelectronics from the University of Strathclyde, Scotland, and a PhD in physics from Heriot-Watt University, Scotland. McArdle is a member of the Institute of Physics, the Optical Society of America, and SPIE.

Contact the authors at the Department of Mathematical Engineering and Information Physics, Faculty of Engineering, University of Tokyo, Hongo 7-3-1, Bunkyo-ku, Tokyo 113, Japan; {ishikawa, neil}@k2.t.u-tokyo.ac.jp.