

Firehose Architectures for Free-Space Optically-Interconnected VLSI Circuits

by

Ashok V. Krishnamoorthy

Bell Laboratories
Lucent Technologies
101 Crawfords Corner Road
Holmdel, NJ 07733-3030

and

D. A. B. Miller

Ginzton Laboratory
Stanford University
Stanford, CA 94305-4085

Preprint

G. L. Report No. 5474
November 1996

To be Published in:
Special Issue on Parallel Computing with Optical Interconnects
Journal of Parallel and Distributed Computing

Firehose Architectures for Free-Space Optically-Interconnected VLSI Circuits

Ashok V. Krishnamoorthy and David A. B. Miller*

Bell Laboratories
Lucent Technologies
101 Crawfords Corner Road
Holmdel, NJ 07733-3030

*Ginzton Laboratory
Mail Code 4085
Stanford University
Stanford, CA 94305

Abstract

Free-space optical interconnects will soon be able to provide input/output bandwidths to a VLSI chip in excess of a terabit per second. The successful application of this technology to parallel distributed processing systems depends on the need for high-bandwidth interconnects and the ability of the architecture to sustain a stream of data at these bandwidths to the processing elements. We review examples of computational tasks that require scalable input/output, that is, computations where the I/O bandwidth of a processing element must grow in proportion to its computational bandwidth. We present several classes of optoelectronic architectures that can support a high-bandwidth data firehose, and mention applications in switching, FFT, sorting, matrix-vector processing, database search, and processor-to-memory interconnect.

1. Introduction: The “Firehose” Problem

The tremendous progress in high-performance, very-large-scale integrated circuit (VLSI) technology has made possible the incorporation of several million transistors onto a single silicon chip with on-chip clock rates of 200 Megahertz (MHz). By 2001, the integration density for silicon complementary metal oxide semiconductor (CMOS) logic is expected to be over 13 million transistors, and the projected on-chip clock rate is expected to be 600 MHz [1]. Estimates made by the Semiconductor Industry Association indicate that the number of transistors available for logic chips and memory chips will respectively double and quadruple every three years. This trend results from two factors: the shrinking feature size of silicon VLSI, resulting in a higher density of gates per unit area, and the improving yield of integrated circuits resulting in more silicon real-estate per chip. The enormous bandwidth that will be available for computation and switching on a silicon integrated circuit will create an increasing demand for high-bandwidth input and output (I/O) to a VLSI circuit. It is widely believed that novel interconnect technologies will be needed to meet this challenge.

One possible solution is the use of three-dimensional optical interconnect technologies via surface normal optical transmitters and receivers. Technologies now exist for attaching GaAs

multiple quantum well (MQW) self-electro-optic-effect device (SEED) photodetectors and light modulators onto a prefabricated silicon integrated circuit using a hybrid flip-chip bonding technique followed by substrate removal of the GaAs chip. This allows surface-normal operation of the optical modulators and detectors [2]. The technique has been used to fabricate high-density, optically-interconnected submicron CMOS integrated circuits by bonding directly above active silicon gates [3]. The procedure effectively decouples the design of the silicon from the placement and bonding of the optical I/O, providing considerable flexibility to the architect of the optoelectronic system. Several thousand operational optical devices on a single chip and digital transmission of data at over 1 Gb/s per device-pair in 0.8 μm CMOS technology [4] have been demonstrated. These results suggest that free-space optical interconnect technologies could soon provide a terabit per second of optical input/output to a conventional silicon VLSI integrated circuit. In contrast, a high-performance interconnection bus can today deliver approximately 5 Gb/s data to a chip [5]; though care must be taken in comparing a future optical technology against a current electrical one (that includes routing and control overhead), it is by no means clear that the electrical technology can be scaled to Tb/s or higher capacities [6]. The emergence of this integration technology, along with free-space and guided-wave (fiber) optical technologies for steering and focusing light beams, and fiber-to-free-space interface technologies, brings new opportunities and challenges to the system designer. One opportunity is that the I/O bandwidth ($\# \text{optical I/O} \cdot \text{transceiver clock-frequency}$) for the processing elements (PEs) can grow *in proportion* to their computational bandwidth ($\# \text{gates} \cdot \text{logic clock-frequency}$). One can approximate the growth of the computational bandwidth by the product of the number of gates and the clock speed of the chip. Even if one makes the conservative assumptions that each pair of optical devices will only function at the clock frequency of the CMOS logic, and that the power dissipation of the chip is governed by SIA predictions, one finds that the total I/O bandwidth of the optoelectronic-VLSI (OE-VLSI) chip scales with the total computational bandwidth of the chip [7]. Furthermore, optical interconnects based on data transmission in fibers offer virtually unlimited distance-capacity products, in stark contrast also to electrical connections whose capacity is strongly dependent on distance. Specifically, for example, the density of information for cabling over meter distances between cabinets or shelves could be increased by approximately three orders of magnitude from an electrical limit of a few Gb/s/cm² (for connectorized coaxial cable) [8].

The question posed in this paper is: “what families of systems could exploit this optical interconnect bandwidth?” We will address this question in two parts. We will first identify computations that require scalable I/O. This subject has been widely studied in the parallel processing literature, and we will cite examples of computational tasks that require PEs to have I/O bandwidths proportional to their computation bandwidth. Having identified tasks that will require vast amounts of I/O bandwidth, we will then address the physical issue of generating and communicating the data at a large aggregate bandwidth to the processing elements. The problem of interfacing conventional, low I/O rate, electronic systems to the very high aggregate rates of parallel optical systems is reminiscent of the notion of trying to drink water from a high pressure firehose. We define a “firehose” architecture as one that can provide sufficient data to balance the optical I/O bandwidth of the processing elements with their computational bandwidth, while still interfacing effectively and efficiently to electronic systems that otherwise have normal electrical interconnects. An important feature of a workable firehose architecture is that the optical I/O bandwidth to the OE-VLSI chip is effectively utilized, while requiring a significantly lower electrical I/O bandwidth to the same chip. Architectures that are unable to provide input data at these high rates, or make excessive demands on the electrical I/O requirements from the chip, become

inefficient in their use of the optical interconnect bandwidth and suffer from a “firehose problem.” This can become a critical issue in exploiting free-space optics. In fact, attempts at building massively parallel optoelectronic computing systems have, historically, had limited success for want of a two-dimensional spatial-light-modulator array that could provide access (via electrical addressing) to the potentially high-bandwidth optical channels [9]; this is simply a restatement of the firehose problem. Consequently, a requirement for firehose architectures is that they possess the means to generate the requisite optical data firehose from external electronic sources that have much lower data bandwidths.

In this paper, we will provide examples of optoelectronic architectures that can solve this firehose problem, and we will classify them by the method they use in creating the high-bandwidth data firehose. In Section 2 we will review the concept of a balanced architecture with a few examples. In Section 3, we will discuss several firehose architectures: (i) the external firehose, (ii) the internally-generated firehose, and (iii) the internal-recirculating firehose. We will also comment on the potential for a “bursty” external processor-to-memory firehose. A brief summary and conclusions will constitute Section 4.

2. The Need for Scalable I/O

In this section, we will review the conditions necessary to balance the computation and communication capabilities of a processing “element” (PE) for a specific task. The purpose is to identify classes of operations that would most benefit from the ability to grow the communication bandwidth of the PEs in proportion to their computational bandwidth. We will conclude that parallel optical interconnects can be used to provide a balanced system for I/O-bound, and certain types of compute-bound, tasks.

To discuss the notion of a balanced system, we use the formalism from Kung [10]. He defines a balanced architecture as one where the computing times and IO times for the PEs are equal for the specific task. Defining the number of operations required in the computation as C_{comp} , and the computation bandwidth of the PE in operations/second as C , then the time required to perform the computations is C_{comp}/C . Similarly, if the number of words the computation requires for I/O is C_{IO} , and the I/O bandwidth in words/second is IO , the communication time is C_{IO}/IO . Hence, for a balanced PE with equal computation and communication times [10]:

$$\frac{C_{comp}}{C} = \frac{C_{IO}}{IO} \quad (1)$$

Thus, when equation 1 holds, neither the processor nor the I/O subsystem is forced to wait for data, and the throughput of the system is maximized. In this paper we will emphasize the technology-independent scaling behavior of this equation in terms of the problem size. We will treat Eq. (1) in the informal sense of equality of scaling behavior, neglecting fixed numerical pre-factors.

As VLSI technology progresses, the computation bandwidth, C , available per PE will increase; however, in general, such improved VLSI technology does not substantially improve the I/O bandwidth of the circuits. Hence, the balance represented in (1) is upset, leading to poor utilization of the VLSI unless we can change the system design in some way. A response to such an

improved technology is to ask each PE to take on more work by increasing the size of the problem given to each PE. Such an approach may or may not work, depending on the relative scaling of C_{comp} and $C_{I/O}$ as the size of the problem is increased.

The most obvious difficulty will occur for so-called "I/O bound" computations, which are computations in which the need for I/O (i.e., $C_{I/O}$) rises at least as fast as the number of computational operations, C_{comp} , as the problem size is increased. Matrix-vector multiplication (with a nonresident matrix) is an example of an I/O-bound computation. $C_{I/O} \sim N^2 + 2N$ I/O operations are required to read in the ($N \times N$) matrix and to handle the (N -dimensional) input and output vectors, and $C_{comp} \sim N^2$ computational operations (multiply-adds) are required. Hence increasing the size (N) of the problem does not rebalance the computation and communication; an increase in C must be matched by an increase in the IO bandwidth, IO , to keep the PE balanced regardless of the size, N , of the problem. Two of the optoelectronic firehose architectures discussed below can be applied to I/O-bound computations.

Tasks where the number of computations required (C_{comp}) rises faster with the problem size N than does the amount of I/O needed ($C_{I/O}$), are referred to as "compute-bound." As the technology improves the number of computations per second (C) that a PE can perform, we can rebalance the PE by having it tackle a larger problem. Hence, at first sight, such "compute-bound" tasks can take full advantage of all such improvements in C . As the PE does more computations for every I/O operation, however, the amount of local memory, M , that the PE needs increases, in general, as some function of $C_{comp}/C_{I/O}$ as the problem size increases [10], with the precise function depending on the mapping of the specific algorithm onto the processors. A problem for such "compute-bound" tasks is that rebalancing the PE without increasing IO may not always be practical because the memory requirements of each PE and the size of the PE may grow too fast (e.g., even exponentially). This can be illustrated by the example of an $N \times N$ butterfly network; this network has applications in multistage interconnection networks (MINs) for switching and sorting, as well as in computing discrete Fourier transforms (DFT). Figure 1 shows an example Butterfly network for the case of $N_{TOT} = 16$ input and output lines. In this example, the individual PEs each have $N = 4$ input and output lines. $\log_N N_{TOT}$ stages of N_{TOT}/N PEs per stage are used to perform the interconnection and/or computations. Figure 1 also shows how a 4-input PE can be made up out of 2-input PEs. In a MIN, each PE performs a cross-connect function on its inputs, depending on the requested destination of the input data. A recursive decomposition of an N_{TOT} -point DFT allows each PE to be built using an $N \times N$ butterfly [11]. In such a DFT-engine, each PE performs an N -point complex FFT on its inputs, and requires $N \log N$ computational operations to do so. Also, each PE requires storage of coefficients for the multiplications, and it needs one coefficient for each input line, i.e., $M = N$. Hence, each PE requires $\Theta(M)$ (i.e., exactly of order M) data-I/O operations and, since $M = N$, $\Theta(M \log M)$ processing operations. The ratio of C_{comp} to $C_{I/O}$ for the array is $\Theta(\log M)$. If the computation bandwidth C is increased by a factor of β and the I/O bandwidth, IO , stays fixed, then the size of each PE's local memory (and the size of the problem N) must therefore be increased from M to M^β to rebalance the PE. Hence, even though this is a "compute-bound" problem, it may eventually not be possible in practice to rebalance the PE because of the exponential increase in memory required to do the rebalancing. This exponential growth in memory requirements can be avoided by employing an optically-interconnected VLSI technology where the communication bandwidth of a PE scales with its processing bandwidth (i.e., where $C = IO$). Indeed, all compute-bound tasks with a C_{comp} to $C_{I/O}$ ratio of $\log M$ (or lower) will have a need for scalable I/O.

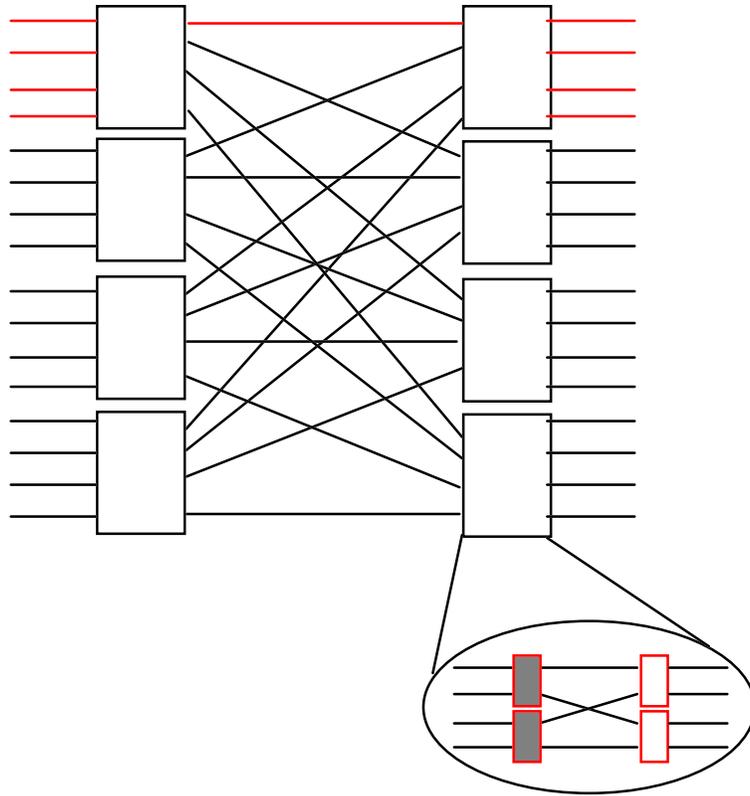


Fig. 1. 16 x 16 multistage interconnection network using two stages of 4 x 4 PEs.

3. Optoelectronic Firehose Architectures

Having identified certain classes of problems that require scalable I/O, we must now design architectures that can exploit the high data bandwidths the optoelectronic PEs are capable of handling. In this section, we will introduce three categories of balanced firehose architectures that can generate and sustain the optical interconnect bandwidth. We will also comment on an unbalanced, bursty, general-purpose, processor-to-memory firehose.

3.1. External Firehose

The first form of data firehose is externally generated. The high-bandwidth data firehose is not created within the optoelectronic system; instead the system relies on some external data generator to supply a high-bandwidth stream of data to the processor elements. We note here that large data firehoses (of 1 Tb/s) do not usually occur in nature. For instance, an uncompressed 10K x 10K pixel natural picture scene (24 bits/pixel) evolving at typical frame rates (30 frames/s) would only generate about 72 Gbits/s data; few natural phenomena evolve at submillisecond time scales or lower. However, data-communications and telecommunications switching applications will demand such bandwidths in the near future. For instance, the tightly-coupled networking of several hundred state-of-the-art processors would be capable of supporting an external-firehose architecture for switching.

In one embodiment of an external-firehose architecture (Fig. 2a), the firehose of data is created incrementally, by combining many individual streams of data into a large data pipe. This method is commonly used in telecommunication systems where a large number (potentially hun-

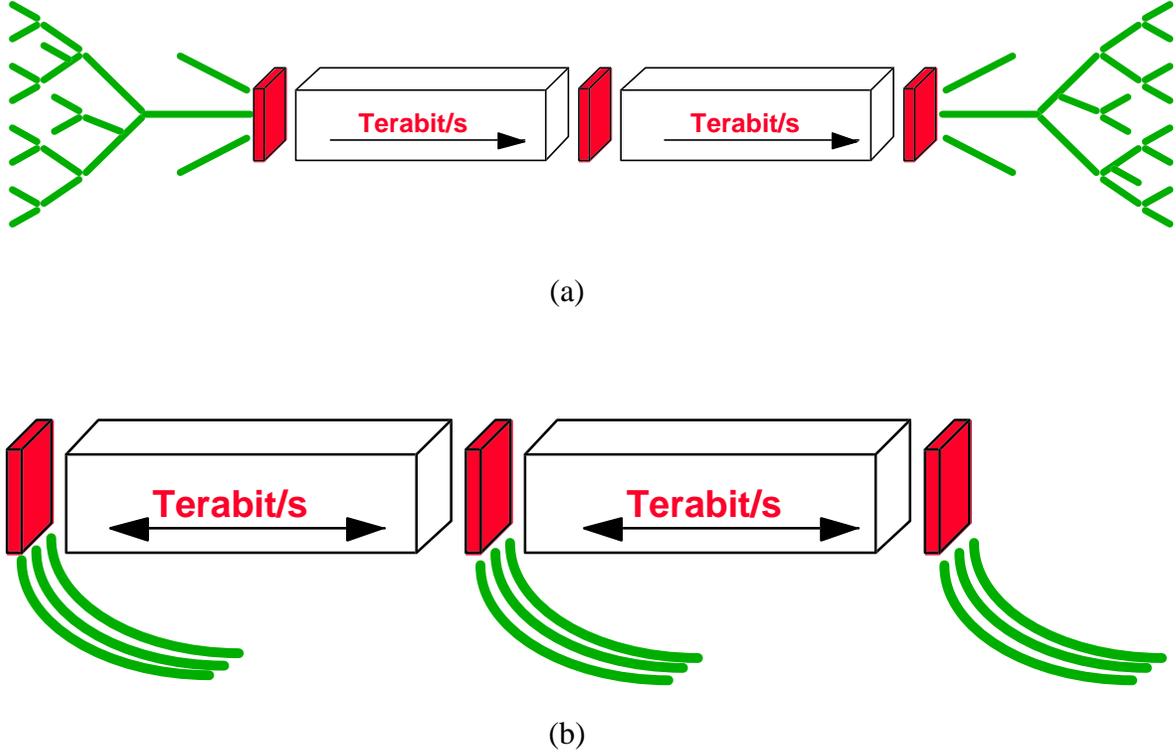


Fig. 2. In an external firehose architecture, the high-bandwidth data stream is an aggregate of many smaller-bandwidth data streams. (a) External firehose with data I/O at the system boundaries. (b) External firehose with distributed data I/O.

dreds of thousands) of smaller-bandwidth data streams are multiplexed into a larger pipe. Free-space optoelectronic switching system prototypes based on this principle have been demonstrated [5,12,13]. In these systems, the aggregate optical firehose is typically delivered to the photonic chips using a two-dimensional fiber bundle. The principle of the external firehose is that an individual input port (or processor) contributes and removes a fraction of the total data that flows through the system. The data generators and extractors may be physically located at the endpoints of the system, or may be distributed throughout the free-space switching fabric in the form of a high-bandwidth bus (Fig. 2b) or hyperplane [14].

We note here that switching is a compute-bound task, even though there is little processing being performed on the data; the primary requirement of the system is to route the data packets. The calculation and the setting of the switches in the network constitutes the processing. When the switch is implemented on a MIN, the calculation of the proper switch states, and the setting of the switches, each require $\Omega(N \log N)$ operations. The lower bound of the ratio C_{comp} to $C_{I/O}$ is then $\Omega(\log N)$, which occurs when the switch is implemented on a MIN with $\Theta(\log N_{TOT})$ stages, as discussed above. Hence, the PE size, N , grows exponentially as C improves if the processor is kept balanced (without increasing I/O), just as for the DFT processor described above. Therefore one can conclude that MINs are well suited to an optoelectronic switching application, given that an external optical data firehose is available to supply the inputs. A similar argument applies to the FFT computation using a multistage network. In fact, the MIN is the most efficient

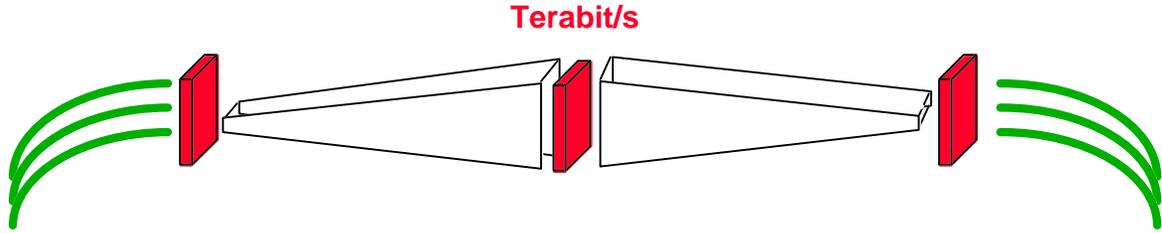


Fig. 3. Internally-generated firehose architecture. The high-bandwidth datastream is created inside the system from a lower-bandwidth input data stream. The high-bandwidth data stream is processed and compressed into a lower-bandwidth output data stream.

architecture for switching in that it minimizes the number of PEs, and hence the processing requirements (but may suffer from internal blocking under certain circumstances). Technological reasons of power dissipation, area, and speed also make photonic implementations of this multi-stage architecture advantageous compared to an all-electronic approach [15]. It should be pointed out that switching may also be performed by a less-efficient (but a strictly nonblocking) architecture such as an $N \times N$ crossbar that has N inputs and outputs and N^2 switches. In this case, the calculation of the switch states again requires at least $\Omega(N \text{Log } N)$ operations, but $\Theta(N^2)$ operations are needed to set the N switches. Thus, an $N \times N$ crossbar requires $\Theta(N)$ I/O operations, and $\Theta(N^2)$ processing operations. Compared to a MIN, a crossbar architecture is easier to rebalance as C increases, but it makes inefficient use of its switching hardware.

3.2. Internally-Generated Firehose

The second form of data firehose is internally generated (Fig. 3). In such a system, the external data bandwidth is modest, but is expanded (in space) inside the system to facilitate processing. For instance, matrix-vector multiplication of a resident $N \times N$ matrix with an N -element vector is an application where the input and output streams have $O(N)$ bandwidth, but the internal bandwidth of the system is $O(N^2)$. If each PE contains an element of the matrix, the (I/O-bound) computation can be achieved by distributing N copies of the i^{th} vector element across the i^{th} row of the matrix, and accumulating the outputs of the PEs along its columns. The use of optical fanout as a means of replicating the $O(N)$ input data solves the interconnect problem and creates an internal $O(N^2)$ data firehose. Algorithms for image processing based on this principle have been suggested [16]. Neural-network architectures also have this data bandwidth requirement, and optoelectronic neural systems based on this principle have been demonstrated [17]. If the matrix is not resident in the PEs, and new matrices must be input to the PEs rapidly, then an additional external firehose for the matrix data is mandated.

3.3. Internal Recirculating Firehose

The third form of data firehose is an internal recirculating firehose (Fig. 4). In this architecture, there is no external source of high-bandwidth data, nor is the input data physically replicated (in space). Instead, the data that is input to the PEs is used multiple times in a system that has a faster aggregate internal data rate compared to its external I/O rate. A fixed amount of information in the form of an optical database recirculates continuously at a high bandwidth onto PEs that use this data as templates. Such a system could, for instance, be used for recognition tasks where the input must be compared rapidly against a fixed, but finite, database (e.g., speech

or image recognition). An example of such a system is a photonic-content, addressable memory system [18]

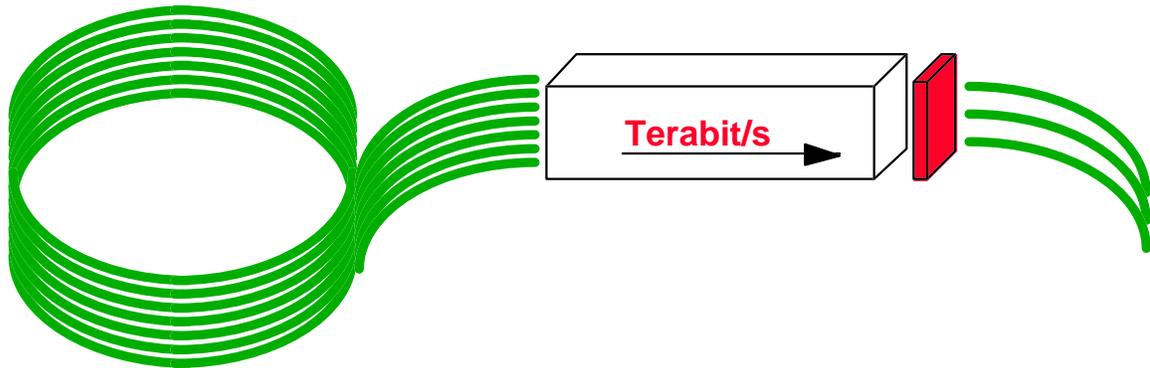


Fig. 4. Internal recirculating firehose architecture. In this architecture, a finite amount of data is continuously recirculated inside the system to create a high-bandwidth data stream. The respective input and output data streams, to and from the system, are relatively small.

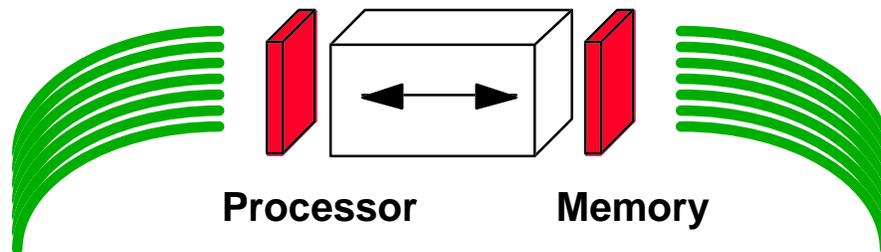


Fig. 5. A processor-to-memory firehose would be an example of an internal bursty firehose architecture, where the internal data stream has a much higher peak bandwidth relative to that of the input and output data streams.

that uses matrix-matrix comparisons followed by summation and thresholding to calculate Hamming distances between the binary input page and a set of stored pages. In this system, the input query to a large database enters the system and is compared at high speed to the recirculating high-bandwidth, optical-database firehose. The results of the high-speed search consist of a number of memory locations; these select memory locations are then accessed and the data is filtered to the output. Due to the simultaneous demands on increasing the information storage capacity and the data bandwidth, this type of database searching is becoming increasingly difficult with conventional means.

3.4. External “Bursty” Firehose

So far we have discussed switching systems and a few application-specific systems that are based on balancing the computation and I/O bandwidths of the PEs. We note that a demand for a high peak bandwidth can also be met with optoelectronic technologies. For instance, future microprocessors may require interconnections to entire chips of cache memory with wide words; this would result in a high-performance, processor-to-memory firehose for general-purpose computing (Fig. 5). In this scenario, low-latency communication would be necessary between the memory chip and the appropriate registers on the processor chip. This would be an external “bursty” firehose; the processing and data I/O bandwidth would not, in general, be balanced, but

the processor would require data (from an optoelectronic cache) to be transferred upon request at a high peak rate.

4. Summary and Conclusions

In this paper we reasoned that free-space, optically-interconnected integrated circuits are a competitive technology for a class of applications that require I/O bandwidths commensurate with their processing bandwidths. In order to make efficient use of the optical interconnect technology, we argued that an optoelectronic architecture must be able to supply a “firehose” of optical data to its processors. This firehose must simultaneously support many parallel channels (1 K) of high-bandwidth data (100 Mb/s). We classified these architectures by the method used in generating the high-bandwidth optical data firehose, providing specific applications for each category of architecture. These include (i) the external firehose, (ii) the internally-generated firehose, and (iii) the internal-recirculating firehose. In the first case, a large volume of high-bandwidth data constitutes the externally-generated firehose. In the second category, a smaller volume of high-bandwidth data is replicated in space to generate the firehose inside the system. Finally, the internal-recirculating firehose uses a fixed quantity of data, but continuously feeds it at a high data rate to the processing elements. I/O-bound problems are obvious candidates for use of optical I/O as processor performance improves. Some important classes of compute-bound problems could also benefit from optical I/O, because otherwise the processor and/or memory size would have to grow too fast as one tried to rebalance the architecture; examples include DFTs and switching. We note also that the criterion for balancing computational and I/O bandwidths is a sufficient, but not necessary, condition for using optical interconnects effectively, and referred to the example of a bursty processor-to-memory firehose with a high instantaneous bandwidth. Finally, we note that the discussion presented in this paper concentrated on data I/O because it was assumed that the data would be loaded optically at high aggregate bandwidths. Future systems may also use optical loading of control signals for MIMD processing. For instance in switching systems, the control signals for an array of switches could be loaded in parallel together with the data. More investigation of a potential “control”-firehose architecture may be warranted.

Acknowledgment: The authors express their gratitude to Fouad Kiamilev, Joseph Ford, Osman Kibar, and Haldun Özaktas for stimulating discussions on the subject of optoelectronic computing architectures.

References

- [1] The National Technology Roadmap for Semiconductors, Semiconductor Industry Association, p. B2 (1994).
- [2] Goossen, K. W., Walker, J. A., D’Asaro, L. A., Hui, S. P., Tseng, B., Leibenguth, R., Kossives, D., Bacon, D. D., Dahringer, D., Chirovsky, L. M. F., Lentine, A. L., and Miller, D. A. B., “GaAs MQW Modulators Integrated with Silicon CMOS,” *IEEE Photonics Technology Lett.* **7** (4), 360-362 (April 1995).
- [3] Krishnamoorthy, A. V., Ford, J. E., Goossen, K. W., Walker, J. A., Lentine, A. L., D’Asaro, L. A., Hui, S. P., Tseng, B., Leibenguth, R., Kossives, D., Dahringer, D., Chirovsky, L. M. F., Kiamilev, F. E., Aplin, G. F., Rozier, R. G., and Miller, D. A. B., “Implementation of a Photonic Page Buffer Based on GaAs MQW Modulators Bonded Directly Over Active Silicon VLSI Circuits,” *Proc. OSA Topical Meeting on Optical Computing*, Salt Lake City, Paper PDP-2 (March 1995).

- [4] Woodward, T. K., Krishnamoorthy A. V., Lentine, A. L., Goossen, K. W., Walker, J. A., Cunningham, J. E., Jan, W. Y., D'Asaro, L. A., Chirovsky, L. M. F., Hui, S. P., Tseng, B., Kossives, D., Dahringer, D., and Leibenguth, R. E., "1Gbit/s Two-Beam Transimpedance Smart-Pixel Optical Receivers," to appear in *IEEE Photonics Technology Lett.* (March 1995).
- [5] Moncayo, A., "Rambus: Bus Design and Analysis at 500 MHz and Beyond," presented at High-Performance System Design Conference (Supercon 1995).
- [6] Miller, D. A. B. and Ozaktas, H., "Limit to the Bit-Rate Capacity of Electrical Interconnects from the Aspect Ratio of the System Architecture," to be Published in *J. Parallel and Distributed Computing*, Special Issue on Optical Interconnects (Submitted October 1996).
- [7] Krishnamoorthy, A. V. and Miller, D. A. B., "Scaling Optoelectronic-VLSI Circuits into the 21st Century: a Technology Roadmap," to appear in *IEEE J. Sel. Top. in Quant. Electron.*, Special Issue on Smart Pixels (1996).
- [8] This limit arises from skin-effect losses due to propagation in ideal coax cables of a few meters in length. See also Ref. [6].
- [9] Neff, J. A., "Perspectives in Optical Computing," *Technical Digest*, OSA Topical Meeting on Optical Computing, Salt Lake City, 176-179 (March 1989).
- [10] Kung, H. T., "Memory Requirements for Balanced Computer Architectures," *J. Complexity* **1**, 147-157 (1985).
- [11] See, for example, Oppenheim, A. V. and Schaffer, R. W., Discrete-Time Signal Processing (Prentice Hall, Englewood Cliffs, New Jersey, 1989), Ch. 9.
- [12] McCormick, F. B., Cloonan, T. J., Lentine, A. L., Saisan, J. M., Morrison, R. L., Beckman, M. G., Walker, S., Wojcik, M., Hinterlong, S. J., Crisci, R., Novotny, R. A., and Hinton, H. S., "Five-Stage Free-Space Optical Switching Network with Field-Effect Transistor Self-Electro-Optic-Effect-Device Smart Pixel Arrays," *Applied Optics* **33** (8), 1601-1618 (March 1994).
- [13] Araki, S., Kajita, M., Kasahara, K., Kubota, K., Kurihara, K., Redmond, I., Schenfeld, E., and Suzaki, T., "Experimental Free-Space Optical Network for Massively Parallel Computers," *Applied Optics* **35** (8), 1269-81 (March 1996).
- [14] Symansky, T. H. and Hinton, H. S., "Design of a Terabit Free-Space Photonic Backplane for Parallel Computing," *Proc. 2nd Intl. Conf. Massively Parallel Processing Using Optical Interconnects (MPPOI'95)*, 16-27 (October 1995).
- [15] Kiamilev, F. E., Marchand, P. J., Krishnamoorthy, A. V., and Esener, S. C., "Performance Comparison Between Optoelectronic and VLSI Multistage Interconnection Networks," *Intl. J. Optoelectronics* **9** (6), 457-470 (September 1994).
- [16] Eshaghian, M. M., "Algorithms for Image Processing on OMC," *IEEE Trans. on Computers*, **40** (7), 827-833 (July 1991).
- [17] Yayla, G., Krishnamoorthy, A. V., Marsden, G. C., and Esener, S. C., "A Prototype 3D Optically Interconnected Neural Network," *Proc. IEEE* **82** (11), 1749-1762 (November 1994).
- [18] Krishnamoorthy, A. V., Marchand, P., Yayla, G., and Esener, S. C., "Photonic Content Addressable Memory System That Uses a Parallel-Readout Optical Disk," *Applied Optics* **34** (32), 7621-7638 (November 1995).