Field-programmable logic devices with optical input–output

Ted H. Szymanski, Martin Saint-Laurent, Victor Tyan, Albert Au, and Boonchuay Supmonchai

A field-programmable logic device (FPLD) with optical I/O is described. FPLD's with optical I/O can have their functionality specified in the field by means of downloading a control-bit stream and can be used in a wide range of applications, such as optical signal processing, optical image processing, and optical interconnects. Our device implements six state-of-the-art dynamically programmable logic arrays (PLA's) on a 2 mm × 2 mm die. The devices were fabricated through the Lucent Technologies–Advanced Research Projects Agency–Consortium for Optical and Optoelectronic Technologies in Computing (Lucent/ARPA/COOP) workshop by use of 0.5-μm complementary metal-oxide semiconductor–self-electro-optic device technology and were delivered in 1998. All devices are fully functional: The electronic data paths have been verified at 200 MHz, and optical tests are pending. The device has been programmed to implement a two-stage optical switching network with six 4 × 4 crossbar switches, which can realize more than 190 × 10^6 unique programmable input–output permutations. The same device scaled to a 2 cm × 2 cm substrate could support as many as 4000 optical I/O and 1 Tbit/s of optical I/O bandwidth and offer fully programmable digital functionality with approximately 110,000 programmable logic gates. The proposed optoelectronic FPLD is also ideally suited to realizing dense, statically reconfigurable crossbar switches. We describe an attractive application area for such devices: a rearrangeable three-stage optical switch for a wide-area-network backbone, switching 1000 traffic streams at the OC-48 data rate and supporting several terabits of traffic.

OCIS codes: 130.0250, 200.4650, 200.3760, 250.3140.

1. Introduction

Conventional electronic integrated circuits (IC's) have a limited electronic I/O bandwidth that is currently constrained to approximately several hundreds of gigabits per second. In contrast, an IC substrate with optical I/O can offer very high optical I/O bandwidths, ranging from several terabits per second (10^{12} bits/s) to potentially hundreds of terabits per second in the future. Hence optoelectronic IC's can alleviate the bandwidth limitations of conventional electronic IC's and can be used in a wide range of applications, such as optical signal processing, optical image processing, and optical interconnects. However, one current limitation to the widespread use of optoelectronic IC's is the need to design a custom silicon VLSI substrate for each different application. This limitation can be addressed by the merging of field-programmable logic with optical I/O, yielding a field-programmable logic device (FPLD) with optical I/O.

FPLD's represent a broad class of programmable devices that encompass both field-programmable gate arrays (FPGA's) and complex programmable logic devices (CPLD's), which differ in only the internal VLSI designs. FPLD's with optical I/O can eliminate the need for the custom design of application-specific optoelectronic VLSI substrates, just as FPGA's have largely eliminated the need for custom design of application-specific VLSI IC's except in the most demanding applications.

Optoelectronic FPLD's can also eliminate the months of turnaround time associated with the fabrication of a custom optoelectronic IC. Currently,

T. H. Szymanski (teds@ece.eng.mcmaster.ca) is with the Communications Research Laboratory, McMaster University, Hamilton, L8S 4K1 Ontario, Canada. M. Saint-Laurent is with Intel Corporation, Hillsboro, Oregon 97123. V. Tyan is with Lucent Technologies, Whippany, New Jersey 07981. A. Au and B. Supmonchai are with the Department of Electrical and Computer Engineering, McGill University, 3480 University Street, Montreal, H3A 2A7 Quebec, Canada.

Received 17 May 1999; revised manuscript received 26 August 1999.

0003-6935/00/050721-12$15.00/0

the design of a custom optoelectronic device can require six months, and fabrication can require several additional months and cost tens of thousands of dollars (depending on the die size). In contrast, the functionality of an optoelectronic FPLD can be programmed dynamically in the field in seconds, typically by the downloading of a control-bit pattern into the device. Optoelectronic FPLD's can also be batch fabricated, leading to a significant cost reduction in comparison with custom devices. Perhaps most importantly, optoelectronic FPLD's can be made compatible with standardized optical I/O pitches, optical packaging assemblies, and optomechanical support structures. The standardization of optical interfaces would seem to be a necessary first step toward the widespread use of optoelectronic IC's.

Integrated field-programmable optoelectronic devices were first proposed in 1994.[1] The proposal was speculative, as no optoelectronic technology was available at that time to support the high gate densities needed to implement field-programmable logic. In 1995 Lucent Technologies announced the development of a complementary metal-oxide semiconductor–self-electro-optic device (CMOS–SEED) technology that merged silicon CMOS with SEED-based optical I/O. A first-generation FPGA with integrated optical I/O was fabricated in the 1995 Lucent Technologies–Advanced Research Projects Agency–Consortium for Optical and Optoelectronic Technologies in Computing (Lucent/ARPA/COOP) workshop, was delivered in 1996 and is described in Refs. 2 and 3. That first-generation device used multiple look-up tables to implement the programmable logic on a 2 mm $\times$ 2 mm die. That device was operational, and several tests that used two optical FPGA devices were performed. That device established the feasibility of integrated optoelectronic FPGA devices, merging tens of thousands of transistors with optical I/O. That device also demonstrated the rapid progress made in the field of optoelectronics from 1994 when logic-gate densities on optoelectronic devices were extremely limited to 1996 when a fully functional FPLD merging tens of thousands of transistors with optical I/O was demonstrated.

There has been recent progress in the development of field-programmable optoelectronic devices. A three-dimensional (3-D) field-programmable system that uses islands of conventional electronic FPGA's arranged in planes interconnected with optical light-emitting diodes between planes was proposed in Ref. 4. A computing system that utilizes programmable optoelectronic processors was proposed in Ref. 5. Field-programmable optoelectronic-device prototypes that use look-up tables were proposed in Refs. 6 and 7, but no experimental results were reported. These designs are similar in concept to our first-generation field-programmable device described in Refs. 1–3. Another field-programmable device was proposed in 1999 in Ref. 8. This proposed device uses holograms to store several programming templates to program the field-programmable logic. It is unclear whether the device was fabricated, and no experimental results were reported. These papers indicate the interest in field-programmable logic with optical I/O.

In this paper, we describe a second-generation optoelectronic FPLD fabricated in the 1997 Lucent/ARPA/COOP workshop. This device demonstrates the Moore law of exponential gains in performance; it is approximately 100 times more powerful than the first-generation field-programmable optoelectronic device described in Refs. 2 and 3. Our latest device uses an array of state-of-the-art programmable logic arrays (PLA's) internally, uses tens of thousands of transistors, and requires approximately 1000 bits of configuration. Our second-generation design is very fast (i.e., hundreds of megahertz), highly functional, and competitive with commercially available FPGA's in performance. Preliminary results on this device were presented in Ref. 9.

Our optoelectronic FPLD designs can be utilized in several ways. It is possible to interconnect several two-dimensional (2-D) optoelectronic FPLD's optically in the third dimension, yielding a 3-D FPLD; see Refs. 4 and 10 for similar concepts. The optical interconnections can be established by use of a fairly standard optical imaging system, as described in this paper. A technology that merges vertical-cavity surface-emitting lasers (VCSEL's) with CMOS substrates is described in Ref. 11. Using lower-wavelength VCSEL technology makes it feasible to stack multiple 2-D optoelectronic FPLD substrates, interleaved with thin passive optical imaging substrates such as diffractive lenslet arrays, directly on top of each other, yielding a compact 3-D FPLD with internal optical interconnections between planes. Three-dimensional VLSI processing systems potentially can allow for the design of previously unattainable high-bandwidth systems in the future.

This paper reviews the architecture, the VLSI design, and the testing of the second-generation device. We also describe an advanced application for these devices: a multiterabit optical packet switch for a wide-area-network (WAN) backbone that uses multiple optically interconnected FPGA's.

2. Architectural Model

Our optoelectronic FPLD consists of a 2-D array of finite-state machines (FSM's) on a CMOS substrate. A single FSM is shown in Fig. 1. Each FSM has several optical and electrical I/O's, a combination PLA, and high-speed memory, i.e., the feedback flip-flops. The digital logic to be performed by each FSM is determined by the PLA-configuration memory, which, in principle, can be loaded electrically or optically. Each FSM has several programmable electronic connections to neighboring FSM cells. For example, a FSM might have connections to its four nearest neighbors in the north, east, south, and west directions. In our optoelectronic FPLD design crossbar switches or additional PLA's can be used to provide programmable interconnections between FSM's (the use of PLA's allows for some additional logic processing to occur between the FSM's.) Each FSM

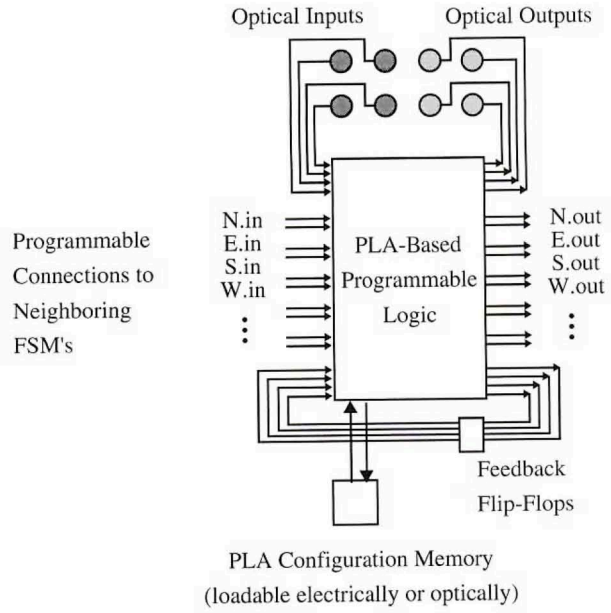

Fig. 1. One FSM implementation of a FPLD combination logic block. Logic functions within the FSM are dynamically programmable. Connections between neighboring FSM's are also dynamically programmable. N, north; S, south; E, east; W, west.

also has several clock and reset inputs that are not shown in Fig. 1.

The core of the FSM is the PLA. Conceptually, a PLA consists of a logical AND array followed by a logical OR array,[12] as shown in Fig. 2(a). The input signals $x_0, \ldots, x_{M-1}$ enter from the lower left-hand side, and the complemented and uncomplemented values of each input signal are fed into the AND array along the vertical lines. Each horizontal line in the AND array represents a product term z_i, where z_i is the logical AND of all the inputs connected to it through a closed connection (represented by the filled circles). For example, in Fig. 2(a) the first product term is $z_0 = x_0$, the second product term is $z_1 = x_1\overline{x_3}$, and the fourth product term is $z_3 = \overline{x_0}\overline{x_1}x_3x_4$. In the OR array each vertical line represents a sum term, i.e., the logical OR of all the product terms connected to it through a closed connection. For example, in Fig. 2(a) $y_2 = z_1 + z_3$. Together these two arrays yield a programmable means of generating any three-level NOT–AND–OR logic function of the input signals.

Conventional PLA's are not dynamically programmable, i.e., the logical function of a conventional PLA is fixed by the pattern of connections [filled circles in Fig. 2(a)] that are permanently embedded in the VLSI array masks at the chip-fabrication time.[12] To make these PLA structures dynamically programmable, it is necessary to insert a programmable cross-point cell at each intersection of a horizontal and a vertical line. Each programmable cross-point cell has two states, opened and closed, and can be controlled by one bit of static RAM (SRAM) control memory. This cell is described in Subsection 2.A. Similar partially dynamically programmable PLA's are used in commercial devices such as the Altera Model MAX 7000 family of programmable logic devices[13] and are often called CPLD's in the literature. Hence our device can also be called a CPLD with optical I/O.

In a general FPLD with optical I/O the programming bits can be entered either electronically or optically. In our current chip, we enter the programming bits electronically from an external FPGA because, for our intended applications, the logic functionality of the device changes relatively infrequently. These programming bits are stored in memory structures called personality matrices (see Subsection 2.A). Within the FPLD the programming bits are distributed to the PLA's on a wide bus and are used to program one column of personality matrices at a time. It is straightforward to adapt the FPLD to accept the programming bits optically, although this approach requires additional hardware to generate and deliver the optical programming bits to the FPLD. In this case an additional FPLD can be user programmed (electronically) to provide optical programming bits, which then can be delivered (by use of fiber ribbons or free space) to program the original FPLD optically. This example highlights the flexibility of FPLD's with optical I/O.

A. NOR–NOR Programmable Logic Array

Figure 2(a) illustrates a PLA that is based on an AND–OR architecture. In practice, a NOR–NOR archi-

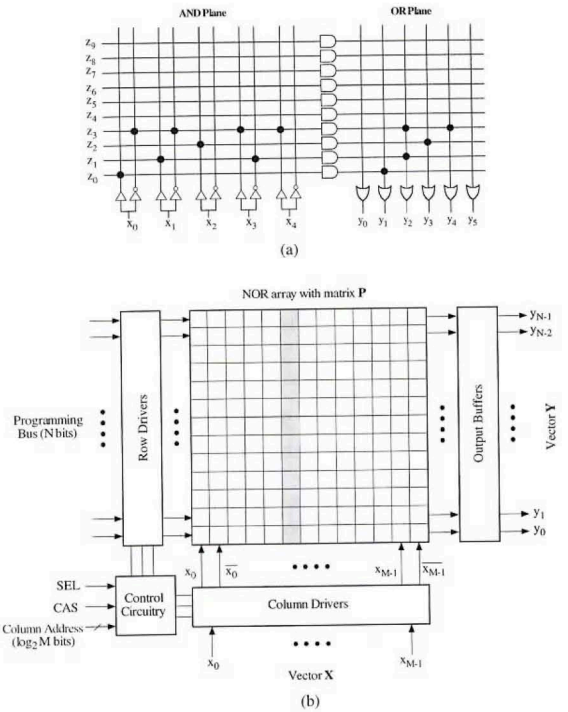

Fig. 2. (a) PLA with an AND array followed by an OR array with vertical inputs $x_0, \ldots, x_{M-1}$, several horizontal product terms $z_0, \ldots, z_9$, and outputs $y_0, \ldots, y_1$. (b) Logic diagram of an individual NOR array with vertical inputs $x_0, \ldots, x_{M-1}$ and with horizontal outputs $y_0, \ldots, y_{N-1}$.

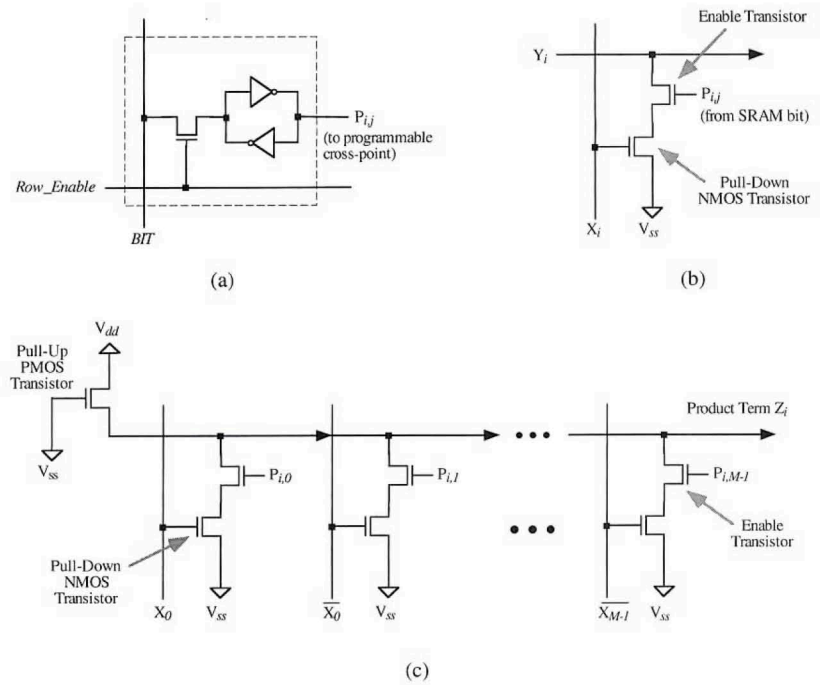

Fig. 3. (a) Five-transistor CMOS SRAM cell used to store one personality bit. (b) Programmable cross-point at every intersection of an input x_i and an output y_i in the NOR array. (c) Dynamically programmable NOR gate with programmable inputs $X_0, \ldots, X_M$ and one output (product) term Z_i. Each NOR gate is a single row of the NOR array shown in Fig. 2(c).

tecture is used in which the AND array shown on the left-hand side of Fig. 2(a) is replaced with a NOR array and the OR array shown on the right-hand side of Fig. 2(a) is replaced with a NOR array. The implementation of an individual NOR array is shown in Fig. 2(b). The input to the NOR array is a vector $\mathbf{X}$ of size M. The output of the NOR array is a vector $\mathbf{Y}$ of size N and is defined as

$$\begin{pmatrix} y_0 \\ \vdots \\ y_{N-1} \end{pmatrix} = \overline{\begin{bmatrix} p_{0,0} & \cdots & p_{0,M-1} \\ \vdots & \ddots & \vdots \\ p_{N-1,0} & \cdots & p_{N-1,M-1} \end{bmatrix} \begin{pmatrix} x_0 \\ \vdots \\ x_{M-1} \end{pmatrix}},$$

$$y_i = \overline{p_{i,0}x_0 + p_{i,1}x_1 + \cdots + p_{i,M-1}x_{M-1}},$$

where

$$p_{i,j} = \begin{cases} 1 & \text{if } y_i \text{ depends on } x_j \\ 0 & \text{otherwise} \end{cases}.$$

The matrix $\mathbf{P}$ can be called the personality matrix; it defines the NOR equations. By cascading two NOR arrays together and redefining the output vector $\mathbf{Y}$ to be the output of the cascaded arrays, we obtain a concise behavioral description:

$$\mathbf{Y} = \overline{\mathbf{P}_2 \cdot \overline{\mathbf{P}_1 \cdot \mathbf{X}}}.$$

Each NOR array is constructed as shown in Fig. 2(b). The NOR array consists of the input vector entering from the bottom, an array of column drivers to create the vertical input signals injected into the center of the NOR array, an array of row drivers on the left-hand side, a 2-D array of programmable cross-point cells programmed by matrix $\mathbf{P}$, and an array of output buffers on the right-hand side. To program the NOR array, we program one column of the personality matrix $\mathbf{P}$ at a time. The control signal SEL is activated to select the NOR array (because there are several NOR arrays on a single chip). The appropriate column-address bits are presented to the control circuitry to select one column of $\mathbf{P}$. The N personality bits are loaded onto a programming bus. Finally, the column-address strobe (CAS) is activated to enable the write operation, which writes the N personality bits from the programming bus to the SRAM in the column of matrix $\mathbf{P}$.

Each programmable cross-point cell consists of a single bit of SRAM, as shown in Fig. 3(a), that controls a two-transistor pull-down circuit, as shown in Fig. 3(b). The single bit of SRAM is based on a modification of the standard five-transistor SRAM cell, as shown in Fig. 3(a). Two CMOS digital inverters are cross-connected to form a storage element. Each inverter requires two metal-oxide–semiconductor (MOS) transistors, for a total of four transistors per stored bit. In addition, a pass transistor is used to allow a bit to be written into (or read from) the storage element. Each programmable cross-point cell therefore requires seven transistors. The geometry of the MOS transistors must be chosen carefully so that a signal activated on the *BIT* line can pass through the pass transistor and remain strong enough to change the value on the storage element. With reference to Fig. 3(a), a write operation is performed by the loading of the data *BIT* line with the

logic value and by the insertion of the control signal $Row\_Enable = 1$.

The two-transistor programmable pull-down circuit is shown in Fig. 3(b). It consists of an enable transistor in series with a pull-down transistor. The personality bit $P_{i,j}$ is connected to the gate of the enable transistor. If $P_{i,j} = 1$, this transistor is enabled and is a closed circuit; otherwise it is an open circuit. If the enable transistor is turned on, the pull-down transistor is connected to the horizontal wire Z_i. The pull-down transistor will pull the value on wire Z_i down to logic 0 (voltage V_{SS}) if its control input is $X_j = 1$. In other words, the logic value on the horizontal wire Z_i is the logical NOR of all the enabled inputs X_i to which it is connected.

A single NOR gate is shown in Fig. 3(c). This NOR gate represents one row of the NOR array shown in Fig. 2(b). There is a single pull-up P-type MOS (PMOS) transistor on the left-hand side, which is always ON and acts as a series resistance to the pull-up voltage V_{dd}. The horizontal wire Z_i passes through an array of horizontal programmable cross-point cells, each comprising a SRAM circuit [as shown in Fig. 3(a)] and a programmable pull-down circuit [as shown in Fig. 3(b)]. The logic value on wire Z_i will be pulled down to logic 0 if any one pull-down transistor is enabled and if its control input X_j is applied. Therefore the circuit shown in Fig. 3(c) implements a programmable NOR gate.

The geometry of the transistors shown in Fig. 3 must be chosen carefully so that the circuit performs properly. It is well known that the transconductance of the PMOS transistors is lower than that of the n-type MOS (NMOS) transistor because of the lower mobility of holes versus electrons. Therefore the PMOS pull-up transistor must be sufficiently large to ensure that Z_i equals logic 1 in the absence of any enabled input $X_j = 1$. Furthermore, the much smaller NMOS transistor must be sufficiently strong that a single transistor can pull Z_i down to logic 0. Typically, the width-to-length ratio (W/L ratio) of the PMOS pull-up is approximately M times the W/L ratio of the NMOS pull-downs, which are of minimum size.

B. Programmable Logic Array Summary

Each of our PLA's that are shown in Fig. 2 contains five inputs, six outputs, and ten product terms. Because each input is complemented and uncomplemented, the NOR array has 10 vertical input signals. Each PLA therefore requires $10 \times 10 = 100$ programmable cross-point cells in the AND array and $10 \times 6 = 60$ programmable cross-point cells in the OR array, for a total of 160 programmable cross-point cells. Each programmable cross-point cell uses the dense five-transistor SRAM cell shown in Fig. 3(a) as well as the programmable pull-down circuit shown in Fig. 3(b). Each FSM also internally buffers its inputs, outputs, and product terms and has a feedback flip-flop, for a total of 22 additional flip-flops that are not shown in Fig. 2.

The programmability of our PLA is computed as follows. Each horizontal line in the AND array can realize a five-input AND gate with as many as five inverters. Each vertical line in the OR array can realize a 10-input OR gate. FPGA complexity usually is expressed in terms of equivalent binary (two-input) logic gates. Each AND-array product term is assumed to be equivalent to six binary gates (a five-input AND gate can be constructed with four binary gates, and five inverters have the same number of transistors as at least two binary gates). A 10-input OR gate is assumed to be equivalent to nine binary gates. Therefore each PLA offers approximately 114 programmable binary logic gates. Each PLA uses 160 bits of configuration memory, for an average of approximately 1.4 bits of control memory per programmable gate. Our entire chip supports six FSM's and requires 960 bits of control memory.

3. VLSI Design

The 1997 Lucent/ARPA/COOP workshop required that all designs be submitted as a 2 mm $\times$ 2 mm die that uses 0.5-μm CMOS technology. With the existing Tanner standard cell library[14] for electronic I/O pads and Lucent Technology standard cells for optical I/O the CMOS substrate supported approximately 200 optical diodes and 40 electronic I/O pins. Hence our device was designed to fit within these parameters.

A picture of one PLA cell is shown in Fig. 4(a). Each PLA requires approximately 300 μm $\times$ 400 μm of VLSI real estate. A photograph of the entire die before the SEED optical I/O were flip-chip bonded onto it is shown in Fig. 4(b). The die also contains some custom smart pixels for an intelligent optical backplane that is under development in Canada by the Canadian Institute for Telecommunications Research (CTIR) and is the subject of a recent paper.[15] We now describe the VLSI design in more depth.

A. Area and Delay of the Programmable Logic Array

Each programmable cross-point cell has dimensions of approximately $24\lambda \times 15\lambda$ for an area of $356\lambda^2$, of which $240\lambda^2$ is used for the SRAM bit.[16] As described in Ref. 17, SRAM densities are approximately $300\lambda^2$–$400\lambda^2$, so our NOR-array design is comparable with state-of-the-art designs. With 0.5-μm CMOS technology each programmable cross-point cell has a width of 18.7 μm and a height of 12.2 μm. The first NOR array [left-hand side of Fig. 4(a)] requires a height of N cells and a width of $2M$ cells.

In this subsection, we formulate the delay of our device analytically and then verify the analysis with SPICE simulations. The RC time constant of a wire of length l, because of the distributed resistance and capacitance of the wire, is approximated as $t_{RC} = (rl)(cl)/2 = rcl^2/2$, where r is the resistance per unit length and c is the capacitance per unit length.

The primary cause of delay in our design is the RC time constant associated with charging and discharging each wire, i.e., $t_{RC} = R_{MOS}C_{Load}$, where the resistance R is determined by the inherent output

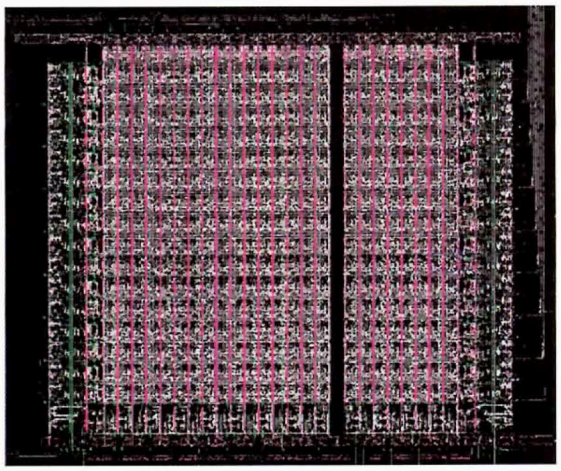

(a)

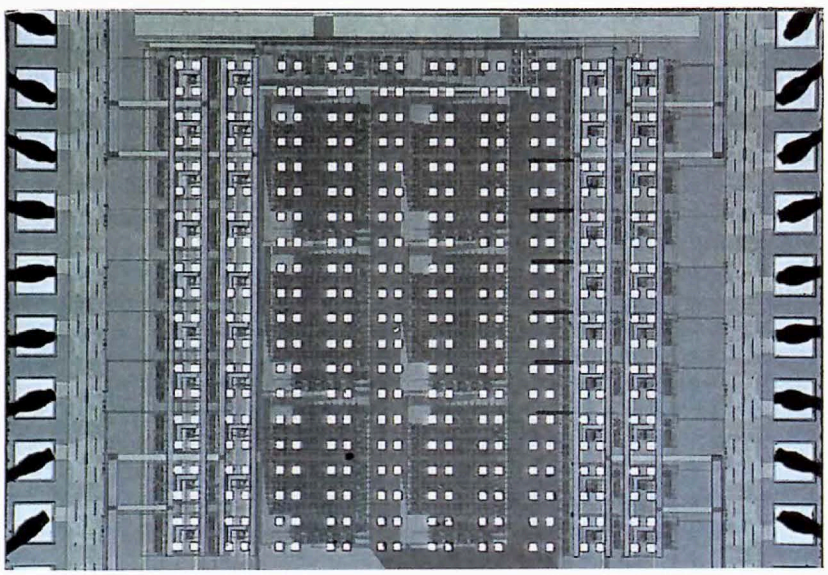

(b)

Fig. 4. (a) VLSI realization of one PLA (300 μm × 400 μm) with five inputs $x_0, \ldots, x_4$, 10 product terms, and six outputs $y_0, \ldots, y_5$. (b) Photograph of the 2 mm × 2 mm die with six PLA's.

resistance of the MOS transistor driving the wire and C represents the total capacitive load. By use of our 0.5-μm CMOS process the capacitance per unit length of metal is 0.056 fF/μm. Therefore the capacitance of a horizontal metal row spanning $2M$ cells is given by

$$\begin{aligned} C_y &= cl \\ &= (0.056\ \text{fF}/\mu\text{m})(18.7\ \mu\text{m/cell})(2M\ \text{cells}) \\ &= 2M(1.05\ \text{fF}). \end{aligned}$$

Some of the vertical control signals in the NOR array were implemented by use of polysilicon because all three metal layers were already used. (See Subsection 3.C for a discussion of this effect.) Polysilicon has a much higher capacitance per unit length than does metal and will result in a larger time constant and slower operation. With our 0.5-μm CMOS process the capacitance per unit length of polysilicon is 0.090 fF/μm. The capacitance of a polysilicon column wire spanning N cells is given by

$$\begin{aligned} C_x &= cl \\ &= (0.090\ \text{fF}/\mu\text{m})(12.2\ \mu\text{m/cell})(N\ \text{cells}) \\ &= N1.10\ \text{fF}. \end{aligned}$$

The above equations establish analytically that the capacitance and the RC time constants of the horizontal and the vertical wires in the PLA are related linearly to the lengths of the wires. In our PLA's the

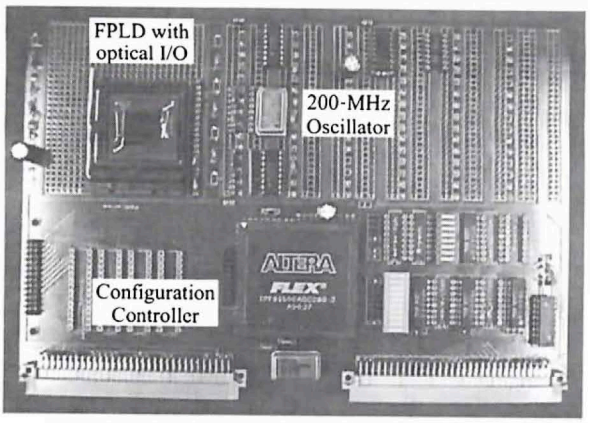

(a)

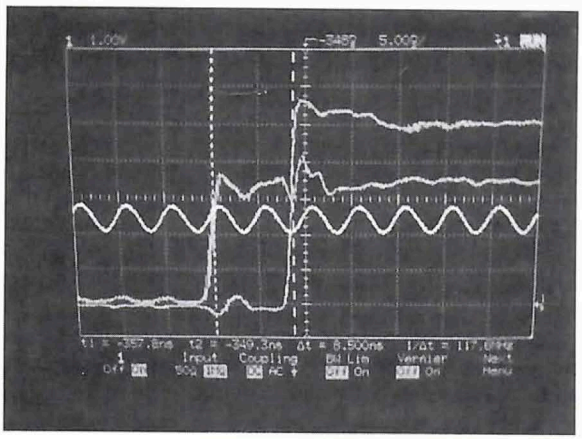

(b)

Fig. 5. (a) PCB used to mount the optoelectronic FPLD for testing. (b) Oscilloscope trace of the test results. The low-amplitude waveform is the 200-MHz optical clock. The first rising edge is the logical input; the second rising edge is the FPLD output.

dimensions are less than 10 cells, so the capacitance of the wires is in the femtofarad range. In addition to the capacitance of the row and the column wires, one must consider the transistor capacitance C_{ds} for all the enable transistors connected to a row wire and the gate capacitance C_{gs} for all the pull-down transistors connected to the column wire. The cumulative capacitance for horizontal and vertical wires with the transistor capacitance taken into consideration is given by

$$C_y = cl + NC_{ds} = c(N \times \text{width}) + NC_{ds},$$

$$C_x = cl + 2MC_{gs} = c(2M \times \text{height}) + 2MC_{gs},$$

respectively, in units of femtofarads.

The above equations verify that the total capacitance is related linearly to the length of the wire. For our 0.5-μm CMOS process the standard capacitive load is approximately 20 fF. We performed detailed SPICE simulations of our NOR gates and NOR arrays and verified the previous analyses, i.e., the total capacitance on the horizontal and the vertical wires is related linearly to their size. We also verified that the total delays on these wires are related linearly to their sizes. From the SPICE simulations our PLA will operate at an approximately 250-MHz clock rate when each wire passes through 10 programmable cross-point cells. The clock rate will increase for smaller PLA's and decrease for larger PLA's.

B. Input–Output: I/O

Four differential optical transceivers provided by Lucent Technologies[18] were implemented on the chip, including the *BL_trnew3_packb*, *BL_trnew3_packa*, *BL_rc5g*, and *BL_trnew3_packc* standard cells. All these circuits are capable of operation at several hundred megahertz clock rates, i.e., 500-MHz clock rates if sufficient optical power is available.[18] The 2 mm × 2 mm die has 200 optical I/O's for an aggregate optical I/O bandwidth of approximately 50 Gbits/s if a clock rate of 250 MHz (the clock rate of our PLA's) is assumed.

The electronic outputs of the chip use the buffered output pads from the Tanner Research standard cell library (*TR_PadOut* cells). According to our SPICE simulations, these electronic pads are rated at 80-, 60-, and 30-MHz clock rates with 20-, 40-, and 80-pF capacitive loads, respectively, which are much slower than that of standard optical I/O cells. According to our SPICE simulations, the Tanner I/O pads have a maximum output clock rate of approximately 30 MHz when driving a standard printed circuit board (PCB) metal trace. The devices were packaged in the standard Kyocera 84-pin pin grid array (PGA) packages used by the Lucent/ARPA/COOP workshop. The electrical output pads are also limited to the perimeter of the die and were restricted to at most 40 I/O pads. Thus the 2 mm × 2 mm die has 40 electrical I/O's for an aggregate electrical I/O bandwidth of approximately 1.2 Gbits/s if a 30-MHz clock rate is assumed. In summary, the optical I/O bandwidth of the device will be considerably greater than the electrical I/O bandwidth, which represents a factor of 40 improvement for this device.

C. Design Methodology

In this subsection, we describe some implications of designing silicon IC's with optical I/O. The inclusion of optical I/O presents several new challenges that are not encountered when designing conventional electronic IC's, which require some changes to the conventional VLSI design methodology.

First, there is a great disparity between the electrical and the optical clock rates, as was established in Subsection 3.B. It may be difficult to design digital logic that can operate at optical clock rates. This effect will be especially important when VCSEL-based optical I/O's are used because VCSEL's can currently operate at approximately 10 GHz and could potentially reach approximately 50 GHz within a few years. Designers of optoelectronic IC's will likely have to change their design methodology to design IC's with several distinct clock domains in which each domain is clocked at a distinct clock rate and with

clock-conversion circuitry between the different domains. For example, the main processing could occur in a slow-speed domain (with a slow-speed clock), and the optical I/O could occur in a high-speed domain (with a high-speed clock), with appropriate conversion circuitry between the domains.

Second, the CMOS IC design *methodology* must be modified when designing with optical I/O by use of standard cell libraries that use three layers of metal (as observed in the design of our first-generation device[3]). In a conventional electronic IC design all three layers of metal are available for the electronic interconnects. Typically, one layer (i.e., metal 1) is reserved for GROUND rails, one layer (i.e., metal 2) is reserved for VDD POWER rails, and the remaining layer (i.e., metal 3) is used for routing signals between standard cells. However, in all optoelectronic technologies that use flip-chip bonding of optical I/O to the VLSI substrate metal 3 must be reserved for the bonding pads of the optical I/O. As a result, designers of optoelectronic IC's may be forced to use polysilicon to route signals between some standard cells, which would result in higher capacitance and slower clock rates when compared with the use of metal wires. This phenomenon explains our use of polysilicon to implement vertical wires, as described in Subsection 3.A.

There are several ways to design around this constraint. For example, computer-aided-design tools can use metal 3 to route signals between standard cells as long as the optical I/O are not nearby, a characteristic that we exploited in our design. However, if the optical I/O are nearby and are using the metal 3 layer for flip-chip bonding, the standard cells underneath the optical I/O must use only two layers of metal and may be forced to use polysilicon.

In general, CMOS technology would be more versatile for optoelectronic processing if an additional layer of metal could be made available for routing signals between standard cells to replace the top metal layer lost as a result of the flip-chip-bonding process. The latest CMOS processes offer as many as six layers of metal, which would simplify the interconnections of the standard cells.

4. Electrical Testing

We completed the electrical testing of our optoelectronic FPLD. A special PCB was designed to test the FPLD, as shown in Fig. 5(a). An Altera Model FLEX 81500 FPGA on the PCB was connected to an IBM PC by use of the parallel port and was used to generate the programming bits for our device. Each test configuration requires specifying and downloading all 960 control bits. We programmed the device to implement several different logic functions. Each programmable cross-point cell in the device was tested and verified to have 100% functionality.

In our most complex test the device was programmed to implement a 12 × 12 switching matrix. The switching matrix includes the first two stages of a three-stage radix-4 Clos network, as shown in Fig. 6. The Clos network has three stages, each with

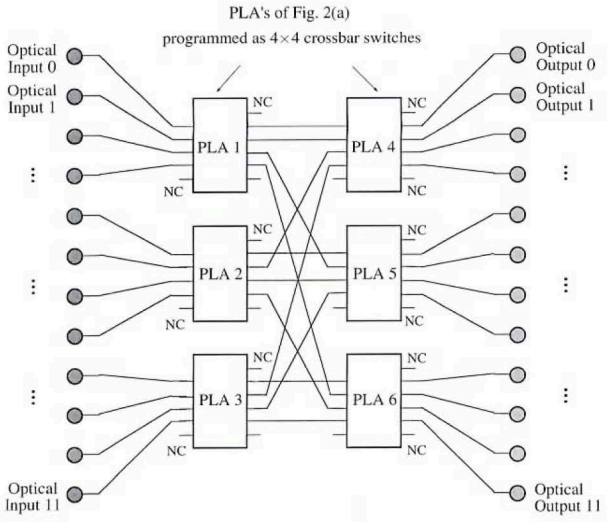

Fig. 6. Optoelectronic FPLD programmed to implement two stages of a three-stage Clos network. NC, not connected.

three 4 × 4 crossbar switches. Each crossbar can be programmed to realize 4! = 24 different programmable permutations. The settings of each crossbar are distinct, and the entire switching network can therefore be programmed to realize $24^6 = 190 \times 10^6$ unique input-to-output permutations. The device can also be programmed to realize all multicasting and broadcasting patterns. This Clos network test was fully functional with a 200-MHz clock rate. In our test, we programmed all the crossbars to realize several mappings, i.e., we could route any input bit stream to any output port. This test illustrates the power of programmable optoelectronic logic devices, as they are competitive with custom-designed optoelectronic devices.[19]

The output waveforms from the oscilloscope for a typical test are shown in Fig. 5(b). The 200-MHz clock has a small voltage swing (2–3 V). One input rising edge is shown and is followed by an output rising edge. A Hewlett-Packard oscilloscope (Model HP 54610B) with two channels and a 500-MHz peak bandwidth was used.

5. Optical Test Setup

We intend to perform the optical testing of the devices by using the setup shown in Fig. 7. We plan to interconnect optically two optoelectronic FPLD's and verify field-programmable digital processing of the optical streams between the devices. Electrical data will be injected into the first device with a 24-channel data generator (Tektronix, Model DG2020). The electrical outputs of the second device will be measured with a logic analyzer (Hewlett-Packard, Model 1661C) and a 500-MHz oscilloscope (Hewlett-Packard, Model 54610B). The devices will be powered optically by an external optical power source, a 100-mW laser diode with an 852-nm wavelength (SDL, Model SDL-5712-H1). A binary phase grating will be used to produce the desired number of

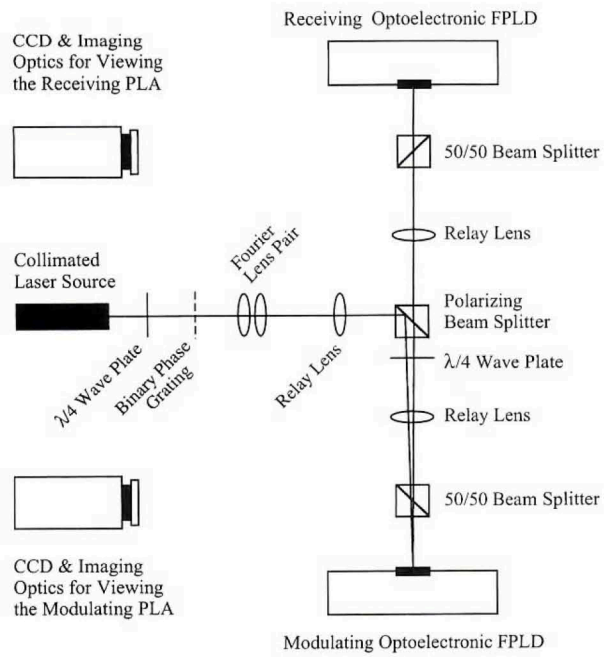

Fig. 7. Optical testing setup.

laser spots on the differential SEED pairs in the PLA. Spindler and Hoyer Model Nanobench optomechanics will be used to construct the optical testing setup shown in Fig. 7. The distance separating the two devices will be approximately 20 cm. This testing apparatus was used successfully to test the first-generation devices that we developed, as described in Refs. 2 and 3.

6. Optoelectronic Packet Switch for a Wide-Area-Network Backbone

Given the rapid growth of data traffic over the Internet the WAN backbone, which links countries and continents together, will soon be expected to support several tens of terabits per second of traffic. Within a decade the WAN backbone may be required to supports hundreds, if not thousands, of terabits per second of traffic. In this section, we describe an attractive application of these devices—an optoelectronic packet switch for a WAN. Although we focus on an individual packet switch in this paper, several of these optical packet switches can be interconnected together in various topologies to yield a scalable WAN with multiterabit capacity.

A. Next-Generation Devices with a 1.4 cm × 1.4 cm Die Size

We first describe the capabilities of our optoelectronic FPLD's, assuming they were scaled to a larger die size of 1.4 cm × 1.4 cm. With 0.5-μm CMOS technology each FSM requires approximately 300 μm × 400 μm of die area. In a 1.4 × 1.4 cm die, we can comfortably fit an approximately 32 × 32 array of these FSM's, leaving 50% of the die area for programmable interconnects and power buses between the FSM's. As was determined from Subsection 2.B, each FSM offers approximately 114 programmable gates, so this chip will offer a peak of approximately 114,000 programmable logic gates. With a modest optical clock rate of 250 MHz the optical I/O bandwidth will be approximately 1 Tbit/s. The device could also use standardized optical I/O packaging and standardized electronic chip-packaging techniques. Hence such a chip could be used in a wide range of applications and could eliminate the need for custom VLSI design of optoelectronic chips.

The use of a CMOS substrate with VCSEL I/O technology promises even higher optical I/O bandwidths and simpler optical interfacing. For example, an optoelectronic FPLD with 1024 VCSEL's, each with a clock rate of 10 Gbits/s, will have an aggregate I/O bandwidth of 10 Tbits/s. It will be difficult for electrical I/O technologies to keep up with these optical I/O bandwidths. CMOS IC's that use electrical I/O pads distributed across the surface of the die can have many more electrical I/O than can conventional perimeter-bound IC's.[20] However, electrical output pads consume moderate to large amounts of power and currently are limited to clock rates of approximately 1 GHz. The newer low-voltage differential–signaling (LVDS) technology allows for electrical I/O with data rates in the few gigahertz range, as demonstrated in the industry. However, these I/O pads require more VLSI area than do regular electronic I/O pads, and they are both larger and slower than state-of-the-art VCSEL I/O technology. It seems that optoelectronic IC's will have much larger optical I/O densities and higher optical I/O clock rates; hence optical I/O bandwidths of hundreds of terabits per second per IC may be available in the future.

B. Wide-Area-Network Backbone

Figure 8 illustrates an optical packet switch for a WAN. A typical long-haul optical fiber cable can contain as many as 12 fiber bundles, each of which can accommodate as many as 12 fibers. Each fiber typically can support as many as 32 bit-serial optical channels by use of wavelength-division multiplexing. In our setting, we assume 12 fiber bundles × 12 fibers × 8 wavelength-division multiplexing channels, of which 1024 channels are used for transmitting optical data. We also assume that each optical channel operates at the Optical Communication-48, or OC-48, data rate standardized by the International Telephony Union at 2.448 Gbits/s.

The optical packet switch shown in Fig. 8 can be constructed by use of a three-stage Clos network, as shown in Fig. 9. Each stage of the Clos network requires 32 crossbar switches, each with 32 inputs and 32 outputs. Hence at the OC-48 data rate this optical switch could simultaneously handle 32 × 32 = 1024 channels clocked at 2.448 Gbits/s, for a total capacity of approximately 2.448 Tbits/s. It may be difficult to maintain a 2.488-GHz clock rate on the entire IC; hence the designer may wish to use two different domains on the IC: one clocked at 2.488 GHz and another clocked at a lower submultiple fre-

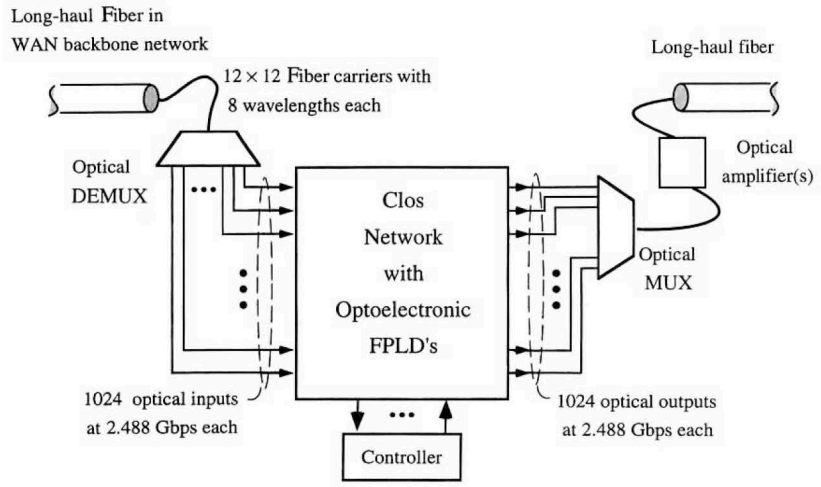

Fig. 8. WAN optoelectronic packet switch based on an optical three-stage Clos network. MUX, multiplexer; DEMUX, demultiplexer; Gbps, gigabits/second.

quency of 622 MHz, as described in Subsection 3.C. The clock rates of the electronic data paths within the switch can be reduced by deserialization and expansion of each bit-serial stream into a 4-bit-wide data path with a clock rate of 2.488 GHz/4 = 622 MHz. This approach will require the use of serial-to-parallel converters at the optical inputs to the IC and parallel-to-serial converters at the optical outputs of the IC. It will also require that the crossbar switches be replicated four times because there will be 4 times as many data streams to switch, although each stream will be at a lower clock rate.

C. Area Estimate

Each stage of the Clos network shown in Fig. 9 can be realized with a single optoelectronic FPLD. The Clos network will therefore require three FPLD's with optical interconnections between the stages. Each FPLD will require 32 crossbar switches of size 32×32 each. Each 32×32 crossbar switch can be implemented by use of a single PLA array, as described in Sections 2 and 3.

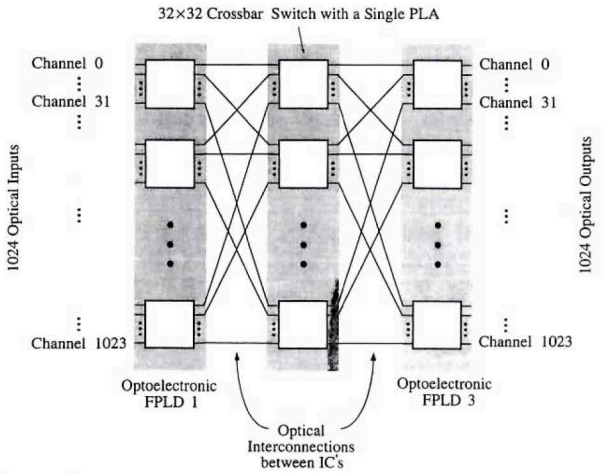

Fig. 9. Optical packet switch with 1024 channels for a WAN that uses a three-stage Clos network that is realized with three optoelectronic FPLD's.

A custom-designed IC could implement the crossbar switches with a single 32×32 NOR array. In other words, for crossbar switching the second NOR array in the PLA of Fig. 2(a) is redundant and can be removed, saving approximately 50% of the area of the PLA. With 0.5-μm CMOS technology a 32×32 NOR array will require a width of (32×18.7) μm and a height of (32×12.2) μm. In addition, the periphery of the array will require an additional 100 μm for the row and the column drivers shown in Fig. 2(b). Therefore the 32×32 crossbar switch will require an area of 700 μm $\times$ 500 μm with 0.5-μm CMOS technology. The drawback of implementing the switches as custom-designed crossbars is that the IC cannot be used for any other logic function because the function is fixed. Therefore the FPLD will implement each crossbar switch as a 32×32 PLA, as shown in Fig. 2(a), with two NOR arrays per PLA and twice the area. In other words, the PLA uses approximately twice as much area as a dedicated crossbar switch with the same number of I/O's but is considerably more flexible because of its programmability.

Each optoelectronic FPLD will require 32 of these switches, or equivalently $32 \times 2 \times (700\ \mu\text{m} \times 500\ \mu\text{m}) = 22.4 \times 10^6\ \mu\text{m}^2$ of silicon area. This area does not include the wiring area between the PLA's. If we include an extra 100% of the area for programmable interconnects and power buses, the total area becomes $44.8 \times 10^6\ \mu\text{m}^2$, corresponding to a die that is roughly 0.67 cm per side with 0.5-μm CMOS technology. This die size is quite small.

Suppose that the optoelectronic FPLD is realized with 0.18-μm CMOS technology. The dimensions required for each 32×32 PLA will decrease by a scaling factor of $0.5/0.18 = 2.8$. Therefore using 0.18-μm technology leads to a single optoelectronic FPGA, thus requiring $32 \times 2 \times (700/2.8\ \mu\text{m} \times 500/$

2.8 μm) = 2.85 × 10^6 $μm^2$ of silicon area. When the area for programmable interconnects and power buses is added (an additional 100%) the total area of the single chip equals 5.7 × 10^6 $μm^2$, or roughly 0.24 cm × 0.24 cm, which is easily realizable.

Without any deserialization logic the IC has only one high-speed clock domain in which all logic must be clocked at 2.488 GHz, i.e., the electronic PLA's must be clocked at the OC-48 data rate of 2.448 GHz. This rate is achievable by careful design of the electronics and by the extensive pipelining of the PLA's. Alternatively, we may design the IC to support two clock domains: a high-speed clock domain at the periphery of the IC and a lower-speed clock domain internally, as discussed in Subsection 3.C. We may add logic to deserialize the bit-serial optical streams at the OC-48 clock rate to a 4-bit-wide stream at an OC-12 clock rate of 622 MHz. This deserialization will require a factor of 4 increase in the hardware because each optoelectronic FPLD must now switch 4096 bit-serial streams at 622 MHz rather than 1024 bit-serial steams at 2.448 GHz. This approach will increase the optoelectronic FPLD dimensions by a factor of 4 to approximately 0.5 cm × 0.5 cm, which is still surprisingly small compared with a current microprocessor die. The entire optical packet switch will require three of these optoelectronic FPLD's.

In the above analyses, we added an additional 100% of the VLSI logic area to account for programmable interconnects and power buses. Our experience is that these elements typically require less than 100% of the VLSI logic area; hence our area estimates should be conservative, leaving sufficient room for interconnect and power. (Furthermore, more advanced processing technologies usually have a larger number of layers of metal available for routing.)

It is possible to implement all three stages of switching crossbars on a single optoelectronic FPLD, although the wiring between the crossbars would be complex and would likely occupy a larger percentage of the die. We chose to partition the design into several smaller FPLD's so that the design would scale upward easily. By partitioning the design across three FPLD's, we can scale our designs upward to optical switches with tens and hundreds of terabits of bandwidth while using only optoelectronic FPLD's with moderate complexity. Finally, we mention that, by the implemention of crossbar switching with optoelectronic FPLD's, the devices can be reprogrammed in seconds to implement any other function as well, resulting in high versatility.

D. Delay Estimate

SPICE simulations indicate that, with 0.5-μm CMOS technology, our PLA's can be clocked at 250 MHz, corresponding to a delay of 4 ns for traversing 10 programmable cross-point cells and settling at the destination latches. With 0.18-μm CMOS technology the delay is reduced by a scaling factor of 0.5/0.18 = 2.8. Thus a clock rate of 2.77/4 ns = 692 MHz is feasible, provided the signals traverse 10 or fewer programmable cross-point cells.

On each IC the data stream must pass through one stage of 32 × 32 crossbars. Each 32 × 32 crossbar can be pipelined, so it takes three clock cycles to traverse the 32 programmable cross-point cells in one NOR array. Hence six clock cycles at the lower clock rate of 622 MHz are required to traverse the crossbar switches in one FPLD. In addition, deserialization requires four clock cycles at the 2.488-GHz clock rate, and serialization requires another four clock cycles at the 2.488-GHz clock rate. Thus the total data delay through a chip is approximately 8 clock cycles at the slower 622 MHz clock rate, which is equivalent to 8(1/622 MHz) = 13 ns/stage of the Clos network.

Recall that the die size with 0.18-μm CMOS technology is approximately 0.5 cm × 0.5 cm. This die is sufficiently small that no significant delays should be introduced by the propagation delays over the wire. The longest wire length is approximately equal to the square root of the chip area, and sufficiently large dies may yield long wires with excessive delays. One can handle excessive delays (i) by adding higher-power buffers to drive the long wires in one clock cycle or (ii) by adding extra pipeline stages and thereby allowing signals to traverse the long wires over several clock cycles. For a die of size 0.5 cm × 0.5 cm wire delays should not be significant and could be handled with either of these techniques.

E. Optical Power Requirements

We now estimate the optical power requirements for the optical FPLD switch first by assuming the SEED technology and then by assuming a VCSEL technology. Assume each SEED optical output is illuminated by a 100-μW beam. If a reflectivity of 50% is assumed along with an optical image-collecting system with a power loss of 3 dB each optical output provides 24 μW of power. This signal must be amplified by use of, for example, an erbium-doped fiber-optic amplifier(s) before injection into the long-haul fiber bundle. If the power-delivery system (including the fiber, the binary phase grating, and the imaging optics) has an attenuation of 6 dB a relatively inexpensive 100-mW laser diode could power 250 optical outputs. Therefore the entire switch with 1024 optical outputs will require approximately 400 mW of optical power, which is moderate and achievable with a few off-the-shelf, relatively low-power laser diodes. SEED's usually operate at a wavelength of 850 nm, which is lower than the wavelengths typically used in long-distance communications, perhaps requiring the use of wavelength converters. The use of VCSEL's may simplify the design because the VCSEL wavelengths can be controlled.

A VCSEL-based FPLD will not require any external optical power supply because the VCSEL's will be electrically powered from the IC. Given 1000 VCSEL's and the assumption of a relatively low current requirement of 0.5 mA each, the current requirements for the optical output will be 0.5 A. If a supply voltage of 3.3 V is assumed the VCSEL's consume ≤2 W of electrical power. The optical receivers and the digital logic will require additional power.

The 1000 receivers may consume an additional 4–8 W, and the digital logic is estimated to consume an additional 20 W (depending on the clock rate). The Semiconductor Industry Association projects the existence within the next decade of VLSI die that can dissipate as much as 200 W of power. Hence the proposed FPLD should be well within the allowable limits of the power dissipation of air-cooled devices. We conclude that, in both designs, the erbium-doped fiber amplifiers will likely require the bulk of the power, followed by the digital logic, followed by the optical I/O.

7. Conclusions

In this paper, we have described the design, the VLSI implementation, and the testing of a second-generation optoelectronic FPLD. We have also described a dynamically programmable PLA structure with a compact VLSI layout and low delays. The NOR arrays of the PLA structure are also ideally suited to implementing statically reconfigurable crossbar switches. We have also described the electrical and the optical testing setups.

The same device scaled to a 2 cm × 2 cm substrate by use of 0.5-μm CMOS technology could support as many as 4000 optical I/O's, could provide 1 Tbit/s of optical I/O bandwidth, and could offer fully programmable digital functionality with approximately 110,000 programmable logic gates. The proposed optoelectronic FPLD is also ideally suited to realizing dense, statically reconfigurable crossbar switches. We have described an attractive application area for such devices: a rearrangeable three-stage optical switch for a WAN backbone that switches 1000 traffic streams at the OC-48 data rate and supports several terabits of traffic.

This device establishes the feasibility of dynamically programmable optoelectronic FPLD's with thousands of transistors per optical I/O bit. Such devices have the potential to reduce significantly the need for the custom design and the fabrication of application-specific optoelectronic devices in the same manner that FPGA's have largely eliminated the need for the custom devices and the fabrication of application-specific gate arrays except in the most demanding applications. These designs can also facilitate the design of 3-D FPLD's by use of optical interconnections between 2-D substrates.

This research was supported by Natural Sciences and Engineering Research Council (NSERC) Canada grant IRG OGP0121601 and by the Canadian Institute for Telecommunications Research (CITR) through the National Centers of Excellence (NCE) program of the Government of Canada.

References

1. T. H. Szymanski and H. S. Hinton, "Architecture of a field programmable smart pixel array," in *Optics in Computing* (Institute of Physics, Bristol, UK, 1995), pp. 497–500.
2. S. S. Sherif, T. H. Szymanski, and H. S. Hinton, "Design and implementation of a field programmable smart pixel array," in *Proceedings of the 1996 IEEE/LEOS Summer Topical Meeting* (Institute of Electrical and Electronics Engineers, New York, 1996), pp. 78–79.
3. S. S. Sherif, S. K. Griebel, A. Au, D. Hui, T. H. Szymanski, and H. S. Hinton, "Field-programmable smart-pixel arrays: design, VLSI implementation, and applications," Appl. Opt. **38,** 838–846 (1999).
4. J. Depreitere, H. Neefe, H. van Marck, J. Van Campenhout, R. Baets, B. Dhoedt, H. Thienpont, and I. Veretennicoff, "An optoelectronic 3D field programmable gate array," in *Proceedings of the Fourth International Workshop on Field Programmable Logic and Applications* (Springer-Verlag, Berlin, 1994), pp. 352–360.
5. M. Ishikawa, "System architecture for integrating optoelectronic computing," Optoelectron. Devices Technol. **9,** 29–36 (1994).
6. D. Fey, B. Kasche, C. Burkert, and O. Tschäche, "Specifications for a reconfigurable optoelectronic VLSI processor suitable for digital signal processing," Appl. Opt. **37,** 284–295 (1998).
7. L. Selavo, S. P. Levitan, and D. M. Chiarulli, "An optically reconfigurable field programmable gate array," in *Digest of the Topical Meeting on Optics in Computing* (Optical Society of America, Washington, D.C., 1999), pp. 146–148.
8. J. Mumbru and D. Psaltis, "Optically programmable gate array," in *Digest of the Topical Meeting on Optics on Computing* (Optical Society of America, Washington, D.C., 1999), pp. 153–155.
9. T. H. Szymanski, M. Saint-Laurent, V. Tyan, A. Au, and B. Supmonchai, "A field programmable gate array with optical I/O," in *Digest of the Topical Meeting on Optics in Computing* (Optical Society of America, Washington, D.C., 1999), pp. 149–152.
10. M. Leeser, W. Meleis, M. Vai, S. Chiricescu, and P. Zavracky, "Rothko: a three-dimensional FPGA," IEEE Design Test. Comput. **15,** 16–22 (1998).
11. A. V. Krishnamoorthy, L. M. F. Chirovsky, W. S. Hobson, R. E. Leibengath, S. P. Hui, G. J. Zydzik, K. W. Goossen, 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 Photon. Technol. Lett. **11,** 128–130 (1999).
12. N. H. E. Weste and K. Eshraghian, *Principles of CMOS VLSI Design: A Systems Perspective*, 2nd ed. (Addison-Wesley, Reading, Mass., 1993), pp. 250–255.
13. *Altera FPGA Data Book 1996* (Altera Corporation, 2610 Orchard Parkway, San Jose, Calif., 1995).
14. *L-Edit Layout Editor Manual*, Version 5 (Tanner Research, Inc., 2650 East Foothill Boulevard, Pasadena, Calif., 1995).
15. T. H. Szymanski and V. Tyan, "Error and flow control in terabit intelligent optical backplanes," J. Select. Top. Quantum Electron. **5,** 339–352 (1999).
16. M. Saint-Laurent, "Software programmable logic array," Undergraduate honors thesis (Department of Electrical and Computer Engineering, McGill University, Montreal, Quebec, Canada, 1997).
17. E. Karali, *Digital Design Principles and Computer Architecture* (Prentice-Hall, New York, 1997).
18. T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, L. A. D'Asaro, L. M. F. Chirovsky, S. Hui, B. Tseng, D. Kossives, D. Dahringer, and R. E. Leibenguth, "1 Gb/s two-beam transimpedance smart-pixel optical receiver mode from hybrid GaAs MQW modulators bonded to 0.8-μm silicon CMOS," IEEE Photon. Technol. Lett. **8,** 422–424 (1996).
19. T. M. Slagle and K. H. Wagner, "Optical smart-pixel-based Clos crossbar switch," Appl. Opt. **36,** 8336–8351 (1997).
20. J. Depreitere, H. van Marck, and J. Van Campenhout, "A quantitative analysis of the benefits of the use of area I/O pads in FPGA's," Microprocess. Microsys. **21,** 89–97 (1997).