A Hybrid-SEED Smart Pixel Array for a Four-Stage Intelligent Optical Backplane Demonstrator

David R. Rolston, David V. Plant, *Member, IEEE,* Ted H. Szymanski, *Member, IEEE,*
H. Scott Hinton, *Senior Member, IEEE,* W. S. Hsiao, Michael H. Ayliffe,
David Kabal, Michael B. Venditti, P. Desai, Ashok V. Krishnamoorthy, *Member, IEEE,*
Keith W. Goossen, *Member, IEEE,* J. A. Walker, B. Tseng, S. P. Hui, J. E. Cunningham, and W. Y. Jan

Abstract—This paper describes the VLSI design, layout, and testing of a Hybrid-SEED smart pixel array for a four-stage intelligent optical backplane. The Hybrid-SEED technology uses CMOS silicon circuitry with GaAs–AlGaAs multiple-quantum-well modulators and detectors. The chip has been designed based on the *HyperPlane* architecture and is composed of four smart pixels which act as a logical 4-bit parallel optical channel. It has the ability to recognize a 4-bit address header, inject electrical data onto the backplane, retransmit optical data, and extract optical data from the backplane. In addition, the smart pixel array can accommodate for optical inversions and bit permutations by appropriate selections of multiplexers. Initial data pertaining to the electrical performance of the chip will be provided and a complete logical description will be given.

I. INTRODUCTION

AN intelligent optical backplane may provide the hardware resources that connection intensive digital systems will require in their ever increasing requirement for bandwidth [1]. By using the two-dimensional imaging properties of free-space optics and the extremely high bandwidth of micron-sized optoelectronic devices, the intelligent optical backplane may far surpass the maximum data throughput of present and future electronic backplanes. Switching systems and distributed processing systems could then be designed to support the ever increasing aggregate bit rates generated by conventional and future silicon processors [2].

The electronics associated with any connection intensive system can at present produce enormous amounts of data, from 1–10 Gb/s per printed circuit board (PCB) [3]. Examples of such systems are video signal processing systems, vision analysis systems, and real-time vector space analysis such as weather prediction models. The bottleneck for most of these types of systems is essentially the accessibility of the resources in the computer system. This is primarily due to the bandwidth and physical limitations of electrical connectors from PCB's. For most multiple PCB architectures, some type of shared resource, such as a bus structure or a sorting node, must be used to move data. With the introduction of a second dimension to connectivity, the former bus structure (with nominally 32 data lines in an electrical backplane) could be extended to a 32×32 array (1024) set of data lines within the same physical space and could be driven at rates comparable with the bandwidth of submicron silicon transistors [4]–[7].

The smart pixel array (SPA) is one method with which to interface the electronic processing capability of silicon with the interconnection capability of free-space-optics [7]. This could be done without requiring a significant change in present-day electronics since the interface to the optical layer would also consist of high-speed silicon chips. The physical layer of the intelligent optical backplane could remain completely transparent to the data processing elements (memory and CPU) where the SPA would simply provide many more interconnection paths for these elements by using the surface area of the chip. Fig. 1 shows conceptually how processors and memories could be interconnected via a dynamically reconfigurable optical backplane.

This paper will describe the electronic design and testing of a simplified SPA. It will begin by briefly describing the interconnection architecture and then give a general overview of the four-board interconnect for which this chip set has been designed. The electrical and optical constraints placed on the design of the chip are then discussed and the chip's logical functionality is given. The electrical and optical testing is then presented and preliminary data is given on the performance of the SPA.

II. SYSTEM DESCRIPTION

An intelligent optical backplane architecture called the *HyperPlane* is being developed. The principles behind the architecture are described in [5] and [7]. The functional specifications for the architecture, smart pixel arrays, and message processors are described in [8]. Before the VLSI design of the smart pixel array was undertaken, the architecture team performed an extensive amount of design and modeling of the architecture, the smart pixel array, and its
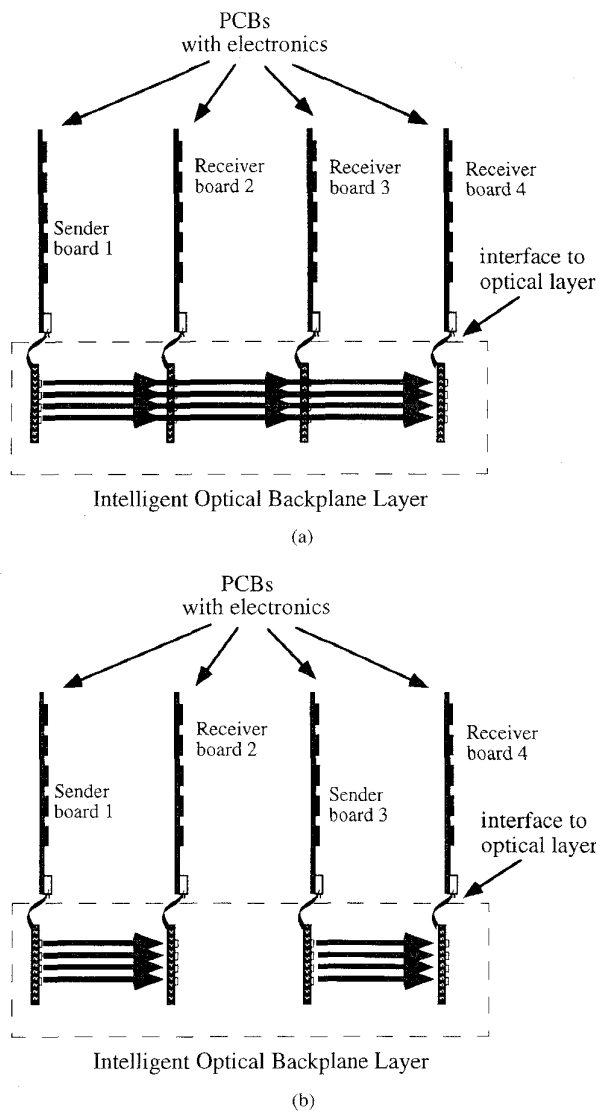
Fig. 1. Conceptual diagram of two interconnections in a reconfigurable optical backplane: (a) broadcast interconnect and (b) mesh interconnect.

associated control electronic circuits at the digital logic gate level [8]. Following accepted system design methodologies, all components were completely specified in the VHDL hardware description language, and all the components were extensively simulated in VHDL to ensure that the demonstrator system would function as intended. The VHDL descriptions were annotated with the best data available for rise times, fall times, logic gate switching delays, and optical propagation delays. After the digital logic design was validated, a VLSI design was undertaken. A more complete description of the digital aspects of the smart pixel array and architecture is found in [7].

The philosophy of this architecture is to provide a large quantity of interconnections among PCB's for packet switching networks and massively parallel processing systems. Specifically, parallel electrical data in packet form is converted into parallel optical data and placed onto specific channels in the optical backplane. The optical data reaches its destination (or destinations) and is then converted back into electrical data. If the packets consist of an address header and data, then a method for address recognition is desirable. An address header can be recognized by a SPA and the subsequent data in the packet can be routed out of the SPA and onto that PCB. If the data is not meant for that PCB, the data is optically retransmitted to the adjacent PCB. In this design, the packet length was chosen to be 32 bits long (four address header bits with 28 data bits). These packets would be sent in parallel 4-bit pieces along a single 4-bit optical channel. Using an externally supplied header clock, the 4-bit address header could be synchronized such that address comparison could take place. By tailoring the address headers of the packets and the addresses of the SPA's with specific patterns, the optical backplane could also be reconfigurable. This would allow the interconnection to implement different interconnection strategies dependent on the application using the backplane or the loads on the network. These interconnects could be versions of fully dilated crossbars, broadcast networks, 2-D and 3-D meshes or other optimal interconnection schemes.

To avoid the latency associated with typical register-based CPU designs and to maintain the fastest possible bit rates, the logical implementation of the SPA must be simple. This would also allow enormous flexibility to the end-users of this technology because only a simple protocol would have to be adhered to in order to use the backplane. Thus, the individual smart pixels are relatively low in logical complexity, composed of only a few multiplexers and address-header recognition circuitry.

This paper analyzes the performance of the Hybrid-SEED SPA and provides a brief overview of the system in which the chip has been designed to operate [9]. A four-board intelligent optical backplane demonstrator was designed to link four SPA's. Each array consisted of four optical channels where a channel was composed of four smart pixels, or a 4×4 SPA. In the smart pixel design described in this paper, only a subset of this 4×4 SPA was fabricated. A single optical channel consisting of four smart pixels was fabricated in a 2×2 pattern in order to be compatible with the optical relay. Thus, the optical relay could be used for both the single channel 2×2 SPA described here and for a 4×4 SPA which will be part of the second iteration of the system.

The optical assembly, including the SPA's, was designed to be placed into a standard electrical backplane chassis. Four external optical power supplies were required to provide arrays of read-out beams to each SPA. A total of 32 read beams are produced per optical power supply and are used to read-out the state of the transmitter of each smart pixel. These read-out beams pass from the transmitters of one stage to the receivers of the next stage with each SPA interrupting the optical path (Fig. 2). The complete system will form a unidirectional ring of four channels. Each of the four SPA's in the backplane were identical and had to be designed with enough flexibility such that they could be interchanged with any node in the backplane. This is of great importance to system designers since it is highly desirable to require only one chip set for all the nodes in the backplane.
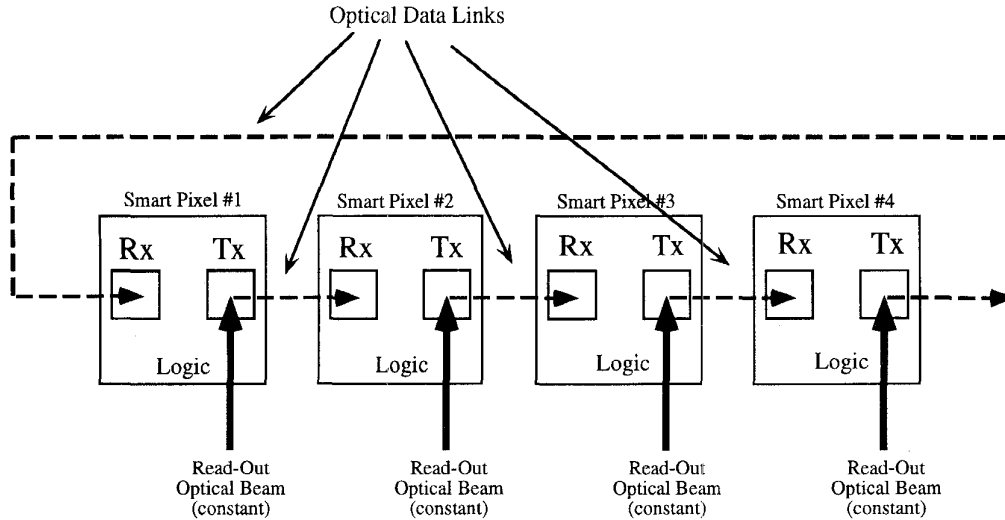
Fig. 2. Conceptual diagram of a four-board modulator-based optical relay showing read-out beams and smart pixel logic.

III. DESIGN CONSTRAINTS

The architecture of the SPA was heavily influenced by the technology of the chip as well as the optical layout. The electrical constraints on design will be covered here followed by several of the optical constraints.

One of the first considerations of the layout was the restriction on the number of electrical bond pads of the chip. The architecture had to allow for a relatively small number of electrical bond pads to service the potentially enormous number of optical smart pixels. In the initial designs, the entire SPA was allowed to have access to and from the electrical bond pads to ease testing, but in future designs, the entire array would not be individually addressable. In a larger SPA, the electrical data lines would have to use a priority scheme to access the backplane. This is analogous to the way a bus master in a typical electrical backplane functions. However, in this case the control of the optical channels would be governed by all the SPA's within the backplane simultaneously.

A second design aspect which did not have to be considered in the design was the clock distribution because the smart pixel was composed of only combinational logic. In this system, there was no need to distribute a clock in order to demonstrate that data could be moved through the backplane. However, this situation is not ideal, because the data is not synchronized as it comes off the backplane. In a larger design, several methods are being explored to distribute clocks to the PCB's within the backplane. These include optical fiber delivery using a pulsed laser diode, or using a designated optical channel within the SPA to broadcast a 50/50 duty cycle square wave. Although this design is unclocked, the architecture is extendible to clocked systems also.

There were several constraints placed on the logical design and physical layout of the SPA due to the optical interconnect designed. The optical interconnect was designed with an image inversion which occurred between adjacent stages, this was done to ensure that the modulated signal beams from one stage were relayed onto the receivers of the subsequent stage. The
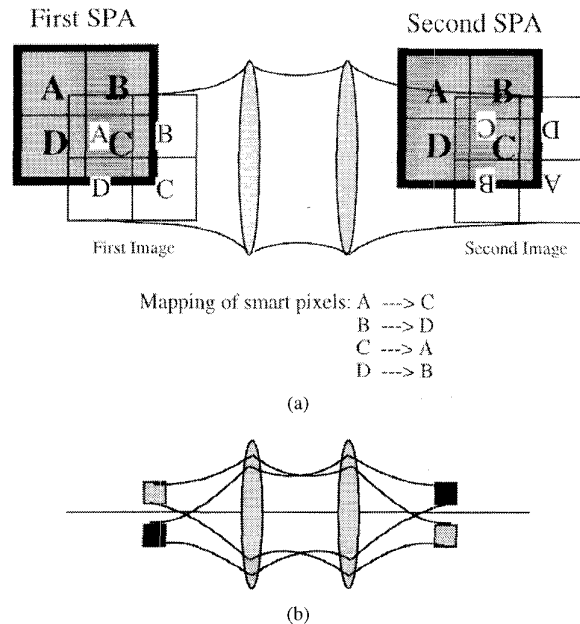

Fig. 3. Conceptual drawing of image rotation: (a) optical bit swapping and (b) optical bit inversion.

array of read-out beams from the transmitters was inverted when it arrived at the receivers. This produced an effect which was called optical bit swapping (Fig. 3), and led to a packet header recognition scheme called 1-hot encoding (to be discussed later). It allowed for a more general and very flexible address recognition circuit to be built. The optical inversion also required that a routing strategy capable of reordering the electrical input and output bits be developed. This was required so that the processing electronics on the PCB's would not be affected by this inversion.

A second logical effect of the interconnect was due to a combination of the optical inversion inherent in the relay and the dual-rail encoding scheme called optical bit inversion. If

a smart pixel sent a logical 1 encoded as a high–low pair of beams, the next smart pixel would receive the optical data as a low–high pair of beams resulting in a logical 0 being received. The optical relay thus acted not only to permute the order of the bits, but also performed a logical inversion of the bits between boards. To negate the effect of this inversion, a multiplexer was incorporated into the receiver of each smart pixel which could select between the received bit and the complement of the received bit.

IV. CONSTRUCTION OF THE SMART PIXEL ARRAY

The silicon chip was a 0.8-μm, three-metal layer, n-well CMOS chip manufactured by Hewlett Packard through the MOSIS foundry service. The chip was then post-processed by Lucent Bell Labs with an array of 20-μm by 60-μm multiple-quantum-well (MQW) modulators and photodetectors using solder-bump-bond technology [10]. Fig. 4(a) is a photograph of the single-channel SPA in the 2×2 configuration. The area outlined shows the position of the single-channel. The inner and outer set of electrical bond pads along the left and bottom of the chip are also shown, with no bond pads along the top and right sides of the chip. This was done to provide access to two independent designs. The outer set of bond pads were specifically for the SPA discussed herein, and the inner perimeter of bond pads were for a design done by the University of Colorado at Boulder [11]. At no time were the inner and outer perimeters of bond pads used simultaneously. Also, the chip as shown is one-quarter of a larger multiproject chip which was diced after the MQW array had been solder-bump bonded, which is why the array of MQW's are only present in the upper-right segment of the chip.

The chip used a three-metal layer process, the highest metal layer (metal 3) was reserved for the MQW diode solder-bump bond pads. Pairs of 20-μm by 20-μm metal 3 pads were placed 20 μm apart for the anode and cathode connections of the MQW diode. The diodes were pitched at 62.5 μm in the horizontal direction and 125 μm in the vertical direction, to form an array of diodes with a 2:1 ratio. The MQW and solder-bump bond is shown in Fig. 5: it consisted of a 90-period stack of AlGaAs and GaAs layers and was tuned for a zero bias exciton peak (λ_0) at a wavelength of 843 nm. A metal reflector was placed below the MQW stack and below that, an ohmic "p$^+$" contact was made. The other side was the connection to the "n" contact using an AlGaAs layer. The MQW diode was used for modulation and detection (0.5 A/W at 850 nm). The characteristics of this device have been well documented and used in several systems [12].

The optical interconnect which supported the SPA was designed such that every second column was used. In Fig. 4(b), the columns of receivers and transmitters are shown; the first and fifth columns consist of receivers, the third and seventh columns consist of transmitters. This led to a uniform 125-μm pitch in both directions for all MQW diodes of the SPA.

The individual smart pixels were designed using four MQW diodes, two for the optical input, and two for the optical output. They function using a dual-rail optical encoding technique so that logical data is encoded. Dual-rail encoding was chosen
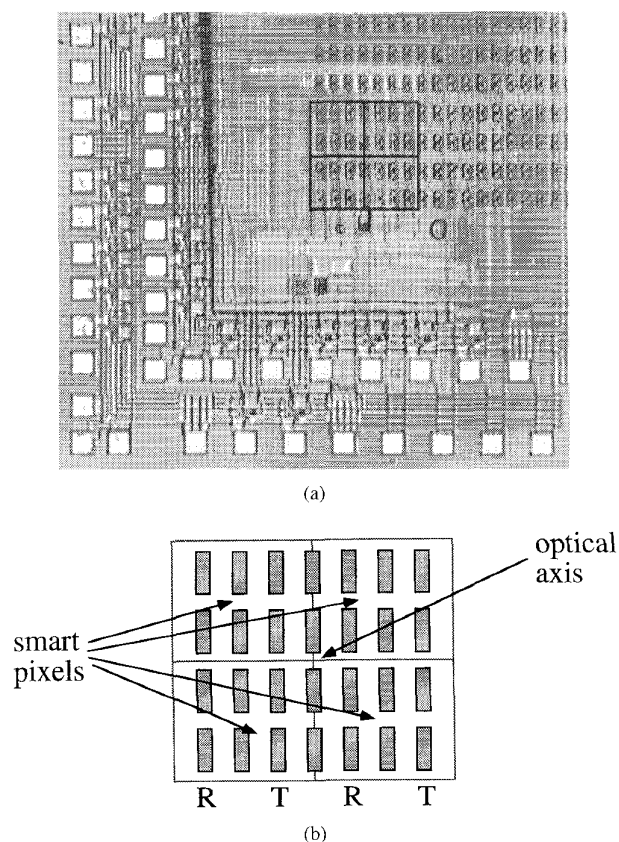

(a)

(b)

Fig. 4.   (a) Photograph of the "2 $\times$ 2" single channel SPA and (b) close-up of (a) showing receiver and transmitter columns.
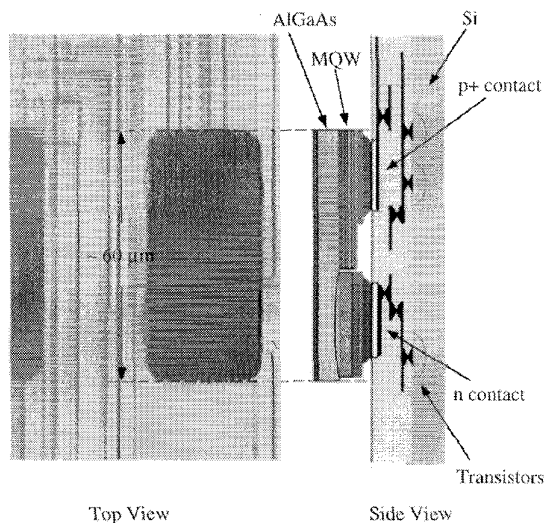

Fig. 5.   Photograph and schematic of solder-bump bonded MQW's.

in order avoid the possibility of level shifting of the absolute optical power between stages. Level shifting can present a serious electrical thresholding problem for receiver designs, especially in these types of cascaded systems with relatively weak signals.

All the MQW diodes were wired in a totem-pole arrangement to provide the dual-rail optical input and output. The
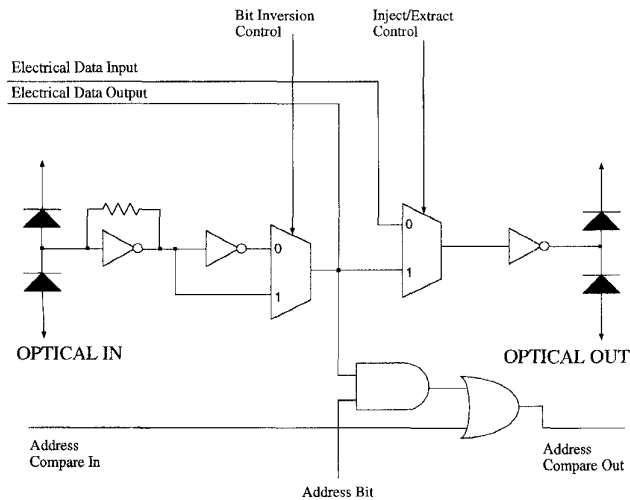
Fig. 6. Logic of a single smart pixel.

Fig. 7. Simplified diagram of the modes of operation of the smart pixel (conceptual): (a) transparent/extract state and (b) inject state.

diodes were biased using four separate supply voltages and allowed for complete flexibility during the testing of the chip. The four bias lines allowed the modulator contrast ratios to be tuned and receiver sensitivity to be maximized. The bias lines consisted of a positive receiver bias, a negative receiver bias, a positive transmitter bias, and a negative transmitter bias. This also allowed for measurements of the photocurrent during the alignment sequence and construction of the optical relay. Prebiasing of the modulators and receivers was considered but was not employed to allow for maximum flexibility in biasing the MQW's.

Each smart pixel consisted of several logical components; a dual-rail optical receiver, two 2-input multiplexers, several inverters, an AND gate, an OR gate, and a dual-rail optical transmitter (Fig. 6). The following is a description of each sub-components within the smart pixel and provides experimental results where possible.

The receiver circuit was a transimpedance amplifier. It used a dual-rail totem-pole MQW pair, a CMOS inverter-amplifier (pMOS width = 12.5 μm and length = 1 μm; nMOS width = 5 μm and length = 1 μm) with two nMOS transistors (width = 2 μm and length = 1 μm) and one pMOS transistor (width = 1.5 μm and length = 1.5 μm) for the feedback path [13]. The receivers were biased using a 5-V power supply. The bias voltages on the receiver MQW's were set at nominally +8.3 V and −3.3 V. The difference in optical power on the totem-pole photodiodes caused current to be either pushed into or pulled out of the receiver, thus causing a change of state in the second inverter. The circuit was simulated with SPICE and showed that a minimum differential optical switching energy of about 20 fJ was attainable using a pair of 20-μm × 20-μm MQW diodes and assuming a sheet capacitance of 0.1 fF/μm^2. This refers to a differential current into and out of the receiver node of 2 μA at a frequency of 100 MHz and would be equivalent to 8 and 4 mW incident optical dual-rail beams. This circuit, and variations of it's basic design have also been used in several experimental demonstrations and has shown that the performance is at least within the same order of magnitude as the SPICE simulations indicate [14].
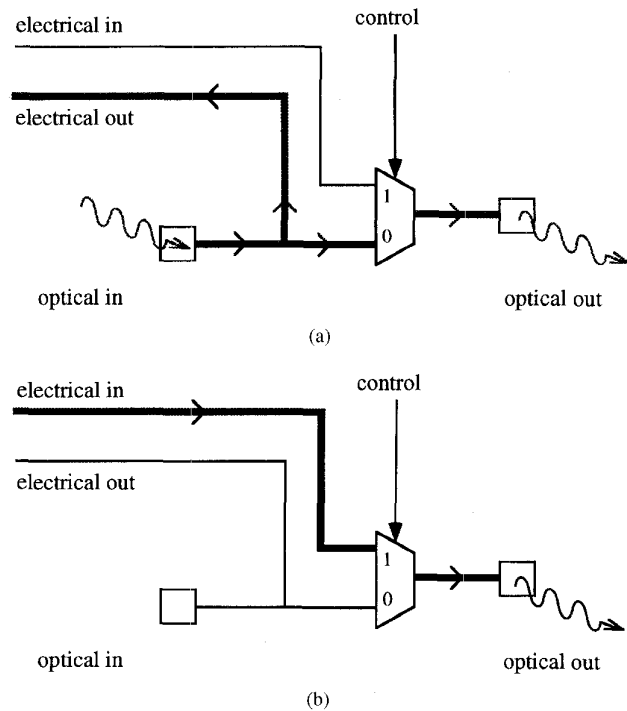
The transmitter circuit was an oversized inverter (pMOS width = 12.5 μm and length = 1 μm; nMOS width = 7.5 μm and length = 1 μm) directly attached to a totem-pole MQW modulator pair [15], and the bias voltages were set at nominally +8.3 and −3.3 V. The capacitance of each diode is on the order of 70–75 fF, including all parasitics such as solder bump [14]. The contrast ratio of the transmitter pair was measured to be 2:1, and the absolute reflectivity was measured to be approximately 30% and 15% dependent on the state.

V. LOGICAL DESCRIPTION OF THE SMART PIXEL ARRAY

The first multiplexer allowed the smart pixel to receive either the direct optical input or the complemented value which was discussed earlier in regard to optical bit inversions. The second multiplexer in the smart pixel was used to select the mode of operation. The two diagrams shown in Fig. 7 show the inject and transparent/extract modes of operation. Note that the extract and transparent modes of operation are identical except that the address recognition bit would be used in the extraction mode.

The specific address recognition circuit designed for this demonstrator was used for two reasons (Fig. 8). The first was to compensate for the optical bit swapping and the second was to allow for a more generalized address recognition called 1-hot encoding. In this way, a simplified address recognition circuit could be made, the permutation of bits could be permitted, and the ability to broadcast packets to multiple boards was possible. If, for example, the address of board 1 was 0001, board 2 was 0010, board 3 was 0100, and board 4 was 1000, the address circuit would produce a high output if at least one match was made. If board 1 sent the address packet
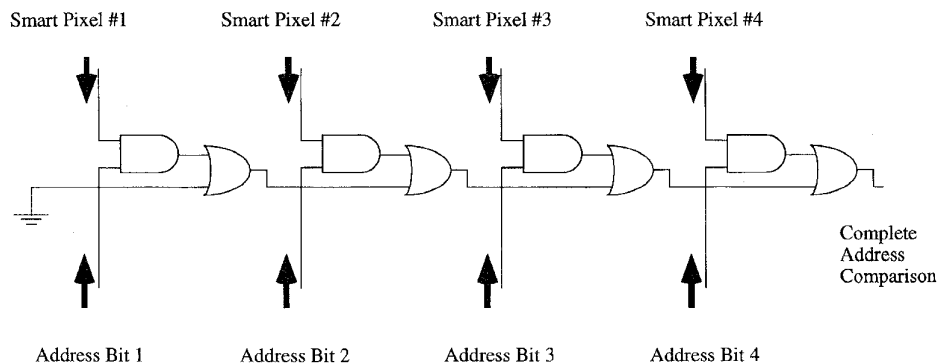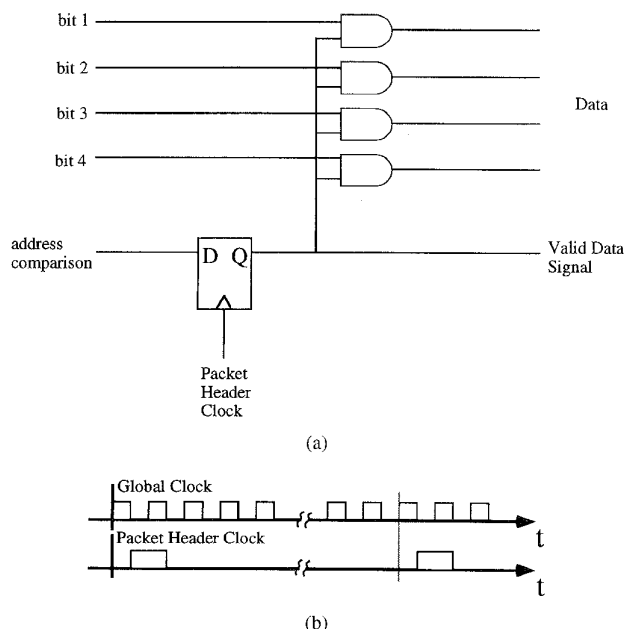
Fig. 8. Address comparison circuit on chip.

(a)

(b)

Fig. 9. External address recognition circuit (proposed).

Fig. 10. 16-bit pattern of all four smart pixels running at 1 Mb/s.
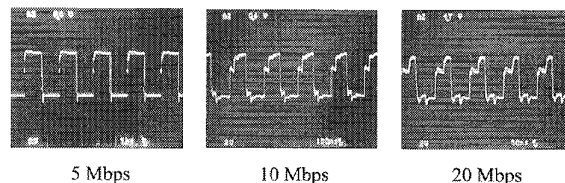
Fig. 11. 5-, 10-, and 20-Mb/s square waves from a single smart pixel.

0101, both board 2 and board 4 would interpret this as a match. Also, since the address would never have more than a single 1 in a definable bit position, the address recognition could be made to work regardless of the optical bit swapping. This method of address comparison has a maximum of five gate delays. This delay is important when longer packet addresses are used, however, other schemes for address recognition are described in [7].

In a fully operating system where address headers and data packets are being used, many schemes could be used in order to synchronize the data stream. However, in order to facilitate the implementation of this chip, the task of incorporating a set of clocks within the SPA design was avoided. External circuitry shown in Fig. 9 could be implemented in order to recognize an address header.

Several strategies were employed to limit the number of electrical off-chip bond pads due to the relatively small chip perimeter. Both the electrical input and output of the smart pixels 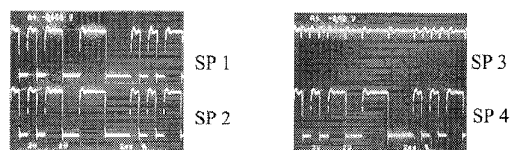could not be routed to their own bond pads, there-fore bidirectional CMOS bond pads were used to multiplex the electrical input and output, where the direction of the bidirectional pads was dependent on the mode of the SPA. The bidirectional pads acted as inputs when the injection mode was selected, and they acted as outputs in the extract and transparent modes. Again, because of the limited number of electrical bond pads, a 9-bit serially loaded configuration register within the SPA chip was used to set-up the state of each smart pixel, where it required only three bond pads to load the register. This register contained static information such as the 4-bit address of the chip, the four bits needed to configure the bit inversion multiplexers, and the bit to control the mode of the SPA (i.e., inject/extract bit).

VI. TESTING OF THE SMART PIXEL ARRAY

The testing procedure was primarily concerned with the logical operation of the SPA. Specifically, the inject, extract, and transparent modes were verified and the configuration of the SPA was tested by loading different patterns into the SPA register. The SPA chip was mounted on a custom printed circuit board (PCB) called a *daughter board*, which supported the SPA chip and a connector. This was done to mechanically isolate the optics and optomechanics from the processor electronics on large 6U PCB's, and thus aid with the stability and alignment of the system [16]. The daughter board

was connected to the processor board via a single high-speed ribbon connector.

Two sets of circuitry were required to operate the SPA, a method for loading the serial register of the SPA and the buffering of the electrical input and output from the bidirectional pads. A finite state machine and parallel-load/serial-out register was built to load the SPA register. The bidirectional electrical I/O were demultiplexed using a set of enabled tristate buffers. The buffers were used to split the bidirectional pads into a set of four input lines and a set of four output lines where the buffers were enabled using the same bit as in the SPA to control inject/extract mode.

The chip was tested using a two-board interconnect using bulk optics and a baseplate. The design of the test-setup was an altered version of a previous demonstration system built, and thus its characteristics and alignability were well known [17]. The optical test rig delivered the eight read-out beams required by the 2×2 SPA with an estimated spot size of 18 μm at a wavelength of 850 nm. In Fig. 10, three of the four smart pixels are shown to be operating with a 1-Mb/s 16-bit pattern; the third smart pixel was not operating because of a slight clipping of its beams (which happen to lie the furthest from the optical axis). Fig. 11 shows a 20-Mb/s square wave pattern which was achieved on a single smart pixel. Although this bit rate was rather low, the speed was limited primarily by the off-chip CMOS pad drivers and it is hoped that future pad driver designs will increase the bit rate. Assuming a 32×32 SPA, with only 32 electrical inputs and 32 electrical outputs, an aggregate bit rate electrically on and off chip of 3.2 Gb/s could be achieved with a 100-Mb/s bit rate per smart pixel. The total optical aggregate bit rate of this SPA (the entire 32×32 array) would be 204.8 Gb/s assuming the 100-Mb/s bit rate per electrical I/O.

The transparent mode was also optically tested. This was a qualitative measurement and could only provide a functional test of the transparent operation of the smart pixel. The bit rates of this mode of operation (optical-in/optical-out) have already been obtained by other researchers who have demonstrated very large data rates [18].

VII. DISCUSSION

The smart pixel array described was used as a first iteration of the optoelectronics in the implementation of a four-stage intelligent optical backplane. This chip was designed to be simple and provide an easy method for the assembly and testing of the optical relay. It provided a large amount of flexibility in the design of the optical relay by allowing for optical bit inversions and optical bit permutations. The complexity of smart pixel arrays will increase in the future and will be capable of handling data synchronization, advanced packet routing, and backplane reconfigurability. Systems will eventually include large numbers of smart pixels, with arrays as large as 32×32 targeted. With these larger systems, significantly more electronic design will be necessary. As described by the general *HyperPlane* architecture, the smart pixel array will be required to act as an interface with the

optical relay where enormous quantities of digital data can be moving at very high speeds.

VIII. CONCLUSION

A smart pixel array containing a single 4-bit-wide optical channel was designed and tested. The SPA was tested both optically and electrically and performed as expected. There were three modes of the SPA: the inject, the extract and the transparent mode. The inject mode allowed electrical data to be placed onto the backplane, and the extract and transparent modes allowed data to be pulled off the chip and retransmitted to adjacent chips in the backplane, respectively. The extraction mode showed a maximum bit rate on and off the chip of 20 Mb/s which was limited primarily by the CMOS output bond pad drivers used.

By demonstrating the three modes of operation of the chip, the basic principles of the *HyperPlane* architecture could be tested. By demonstrating the address recognition of the SPA, a degree of confidence was achieved such that a more aggressive implementation of the chip could be undertaken. The chip also provided a flexible means in which to exercise the optical relay and demonstrate a multiboard optical backplane.

REFERENCES

[1] R. A. Nordin, F. J. Levi, R. N. Nottenburg, J. O'Gorman, T. Tanbun-Ek, and R. A. Logan, "A system perspective on digital interconnection technology," *J. Lightwave Technol.*, vol. 10, pp. 811–827, 1992.

[2] T. Koinuma, and N. Miyaho, "ATM in B-ISDN communication systems and VLSI realization," *IEEE J. Solid-State Circuits*, vol. 30, pp. 341–347, Apr. 1995.

[3] L. A. Lev, A. Charnas, M. Tremblay, A. R. Dalal, B. A. Frederick, C. R. Srivatsa, D. Greenhill, D. L. Wendell, D. D. Pham, E. Anderson, H. K. Hingarh, I. Razzack, J. M. Kaku, K. Shin, M. E. Levitt, M. Allen, P. A. Ferolite, R. L. Bartolotti, R. K. Yu, R. J. Melanson, S. I. Shah, S. Nguyen, S. S. Mitra, V. Reddy, V. Ganesan, and W. J. de Lange, "A 64-b microprocessor with multimedia support," *IEEE J. Solid-State Circuits*, vol. 30, pp. 1227–1235, Nov. 1995.

[4] 1994 SIA Roadmap, "Semiconductor Industry Association," San Jose, CA, 1994.

[5] H. S. Hinton and T. H. Szymanski, "Intelligent optical backplanes," in *Proc. Second Int. Conf. Massively Parallel Proc. Using Opt. Interconnections*, E. Schenfeld, Ed. Los Alamitos, CA: IEEE Comput. Soc. Press, Oct. 1995, Cat. no. 95TB100004, pp. 133–143.

[6] IEEE Standard for Futurebus+, IEEE Standard 896.2, P. M. Kelty, Ed., 1991.

[7] T. H. Szymanski and H. S. Hinton, "Reconfigurable intelligent optical backplane for parallel computing and communications," *Appl. Opt.*, vol. 35, no. 8, pp. 1253–1268, Mar. 1996.

[8] T. H. Szymanski, P. Desia, M. Kim, and S. Gagnon, "Phase II architecture, smart pixel array, message processor, and system testing specifications," Internal CITR report, 1995.

[9] D. V. Plant, B. Robertson, H. S. Hinton, M. H. Ayliffe, G. C. Boisset, D. J. Goodwill, D. Kabal, R. Iyer, Y. S. Liu, D. R. Rolston, M. B. Venditti, T. H. Szymanski, W. M. Robertson, and M. R. Taghizadeh, "Optical, optomechanical, and optoelectronic design and operational testing of a multi-stage optical backplane demonstration system," to be published.

[10] K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.,* vol. 7, pp. 360–3622, Apr. 1995.

[11] K. E. Devenport, H. S. Hinton, D. J. Goodwill, D. V. Plant, D. R. Rolston, and W. Hsiao, "A hyperplane smart pixel array for packet based switching," in *IEEE/LEOS Summer Top. Meet.,* Keystone, CO, Aug. 5–9, 1996, submitted.

[12] A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. T. Tseng, R. E. Leibenguth, D. P. Kossives, D. W. Dahringer, D. D. Bacon, T. K. Woodward, and D. A. B. Miller, "Arrays of optoelectronic switching nodes comprised of flip-chip-bonded MQW modulators and detectors on silicon CMOS circuitry," *IEEE Photon. Technol. Lett.,* vol. 8, pp. 221–223.

[13] 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. P. Hui, B. Tseng, D. Kossives, D. Dahringer, and R. E. Leibenguth, "1-Gb/s two-beam transimpedance smart-pixel optical receivers made from hybrid GaAs MQW modulators bonded to 0.8-mm silicon CMOS," *IEEE Photon. Technol. Lett.,* vol. 8, no. 3, pp. 422–424, Mar. 1996.

[14] A. V. Krishnamoorthy, T. K. Woodward, R. A. Novotny, K. W. Goossen, J. A. Walker, A. L. Lentine, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovski, G. F. Aplin, R. G. Rozier, F. E. Kiamilev, and D. A. B. Miller, "Ring oscillators with optical and electrical readout based on hybrid GaAs MQW modulators bonded to 0.8 μm silicon VLSI circuits," *Electron. Lett.,* vol. 31, no. 22, pp. 1917–1919, Oct. 1995.

[15] D. A. B. Miller, "Novel analog, self-electrooptic-effect devices," *IEEE J. Quantum Electron.,* vol. 29, pp. 678–698, Feb. 1993.

[16] D. V. Plant, B. Robertson, H. S. Hinton, M. H. Ayliffe, G. C. Boisset, W. Hsiao, D. Kabal, N. H. Kim, Y. S. Liu, M. R. Otazo, D. Pavlasek, A. Z. Shang, J. Simmons, K. Song, D. A. Thompson, and W. M. Robertson, "A 4 × 4 VCSEL/MSM optical backplane demonstrator system," *Appl. Opt.,* vol. 35, no. 29, Oct. 1996.

[17] D. V. Plant, A. Z. Shang, M. R. Otazo, D. R. Rolston, B. Robertson, and H. S. Hinton, "Design, modeling, and characterization of FET-SEED smart pixel tranceiver arrays for optical backplanes," *IEEE J. Quantum Electron.,* vol. 32, Aug. 1996.

[18] A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T. K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "3-D integration of MQW modulators over active submicron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter circuit," *IEEE Photon. Technol. Lett.,* vol. 7, pp. 1288–1290, Nov. 1995.

David R. Rolston was born in 1971 in Ottawa, Canada. He received the B.Eng. and M.Eng. degrees from McGill University, Montreal, PQ, Canada, in 1993 and 1996, respectively. He is currently working towards the Ph.D. degree at McGill University.

He has been working on several aspects of the design and development of photonic backplanes involving fee-space optics. He has been working on two aspects of the design of optical baclplanes since beginning the Ph.D. The first is an analaysis of a high-density optical interconnect scheme using clustered window patterning. The second is the design, layout, fabrication, and testing of CMOS smart pixel arrays with solder-bump bonded MQW receivers and modulators developed by Lucent Bell Labs.

David V. Plant (S'86–M'89) was born in 1962 in Vancouver, Canada. He received the B.S., M.S., and Ph.D. degrees from Brown University, RI, in 1985, 1986, and 1989, respectively.

From 1989–1993, he was a Research Engineer in the Department of Electrical Engineering at the University of California at Los Angeles. In 1993, he joined the Department of Electrical Engineering at McGill University, Montreal, PQ, Canada, as an Assistant Professor. His research interests include optoelectronic devices and packaging, optical interconnects, and free-space digital optics.

Dr. Plant is a member of the Optical Society of America and the American Physical Society. He is on the technical program committees for the 1996 IEEE Summer Topical Meeting on Smart Pixels, the 1996 LEOS Annual Meeting, and the IEEE Meeting on Massively Parallel Processing using Optical Interconnects (MPPOI '97).

Ted H. Szymanski (S'84–M'87) received the B.Sc. degree in engineering science and the M.A.Sc. and Ph.D. degrees in electrical engineering from the University of Toronto in 1980, 1982, and 1988, respectively.

From 1982 to 1984, he was consulting for various companies in the Toronto area, working primarily in distributed operating systems. From 1987 to 1991, he was an Assistant Professor at Columbia University, New York, NY, and a Principle Investigator at the NSF Center for Telecommunications Research, working on ATM networks and WDM optical architectures. He is currently an Associate Professor at McGill University, Montreal, Canada, and a Project Leader in the Canadian Institute for Telecommunications Research (CITR), leading a project on "Optical Architectures and Applications." He codeveloped the *HyperPlane* multichannel intelligent optical backplane architecture currently being pursued by the CITR. His research interests include ATM switching and computing architectures, performance analysis and congestion control of ATM networks, and fiber-based and free-space-based optical architectures.

H. Scott Hinton (S'81–M'82–SM'92), for photograph and biography, see this issue, p. 23.

W. S. Hsiao, photograph and biography not available at the time of publication.

Michael H. Ayliffe was born on August 8, 1973, in Lachine, Canada. He received the bachelor's honours degree in electrical engineering from McGill University, Montreal, Quebec, Canada, in 1996. He is currently working towards the M.S. degree at McGill University.

In the past two years, he has been working on various aspects of the design and development of free-space optical backplanes. His major field of concentration has been the design and development of optomechanical hardware for free-space optical interconnects. Other research interests include smart pixel arrays as well as the analysis, design, and characterization of optical interconnections.

David Kabal was born on February 5, 1972, in Toronto, Canada. He received the bachelor's degree in electrical engineering from McGill University, Montreal, Quebec, Canada, in 1995. He is currently working towards the M.S. degree in electrical engineering at McGill University. His thesis research is concerned with the design of electrical and optomechanical packaging for two-dimensional smart pixel arrays.

Michael B. Venditti was born on November 13, 1974, in Montreal, Quebec, Canada. He is currently working towards the B.Eng. degree at McGill University, Montreal, Quebec, Canada.

His academic and research interests include optics, optoelectronic devices and packaging, communications systems, and digital electronics.

P. Desai, photograph and biography not available at the time of publication.

Ashok V. Krishnamoorthy (S'90–M'93), for photograph and biography, see this issue, p. 76.

Keith W. Goossen (S'87–M'88), for biography, see this issue, p. 83.

J. A. Walker, photograph and biography not available at the time of publication.

J. E. Cunningham, photograph and biography not available at the time of publication.

B. Tseng, photograph and biography not available at the time of publication.

W. Y. Jan, photograph and biography not available at the time of publication.

S. P. Hui, photograph and biography not available at the time of publication.