PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Photonic Switching Technology [and Discussion]

J. E. Midwinter and W. J. Stewart

Phil. Trans. R. Soc. Lond. A 1989 **329**, 119-130
doi: 10.1098/rsta.1989.0063

Phil. Trans. R. Soc. Lond. A **329**, 119–130 (1989) [119]
Printed in Great Britain

Photonic switching technology

By J. E. Midwinter, F.Eng., F.R.S.
University College London, Torrington Place, London WC1E 7JE, U.K.

I examine the scope for exploiting the properties of free-space and guided-wave interconnects to enhance the performance of synchronous self-routing switching matrices, with application to fast packet or asynchronous time division switching in mind.

1. Introduction and network requirement

The success of optical fibres in telecommunications and other forms of wideband transmission is having the immediate effect of removing the long-standing bottleneck in transmission within telecommunication networks. This has shifted the emphasis to the problems of switching and routing signals and has done so at a time when there is growing interest in carrying a very wide variety of signals spanning a huge range of call-set-up times and transmission data rates, e.g. a message length of a few hundred bits in length, to database file transfer involving many gigabits per second of data and at rates from 64 kbit s^{-1} telephony to cable television (CATV) or high-definition TV at 100–600 Mbit s^{-1}. This can result in total data flows per connection spanning 10^3 to 10^{13} bits, an astonishing range of 10 orders of magnitude!

To handle such a range, some form of 'bandwidth on demand' assignment is required and a common theme among many proposals to achieve this, notably the European Community's (EC's) Research on Advanced Communications for Europe (RACE) programme, is the asynchronous time division (ATD) data format whereby a wideband channel is broken up into a very large number of rigidly formatted and regularly spaced data packets, each of which can be individually addressed. As a result, bandwidth assignment is reduced to packet assignment and may involve 'posting' one single packet or hundreds or thousands per second, according to need. At a major switching node, this can clearly generate a requirement for switches of awesome capability.

2. Capability of optics

Noting the clear message from optical communications that optics brought to transmission the attributes of low attenuation, low cross talk and wide bandwidth, we look to see how these might assist switching. Concentrating on the fibre transmission bandwidth or spectrum, it can be used as a guided-wave 'free space' in which signals transmitted through it are coded for their destination by the optical carrier frequency or wavelength. Such an approach is commonly described as dense wavelength division multiplexing (DWDM) and may support 10–50 channels by using incoherent detection (Bulley *et al.* 1987) or many more by using coherent detection (Stanley *et al.* 1987). It is discussed elsewhere in this meeting.

The second approach is to copy radar technology and to code each transmitted pulse as a pulse sequence, typically of pseudo-random form or to use some form of spread spectrum

coding (code division multiple access or CDMA). In this way, a suitable matched filter placed in front of the optical receiver preferentially selects the data destined for itself and discriminates against all other data. However, with present incoherent detection technology, this route does not look too promising since the required bandwidth appears to extrapolate roughly as the cube of the number of data channels (Chung *et al.* 1989). Both the above techniques code data at source, so that it self-routes through a passive network. However, given that very large data rates can be handled by simple electronic circuits (say to 30 Gb s^{-1}) but that the timing and communication problems escalate rapidly as the circuit complexity and clock rate increase, I examine here the idea that optics may be used as a sophisticated interconnect to assist either electronic or optical logic towards higher clock rate operation.

3. All-optical logic and optoelectronic interface devices

Optical components are capable of operating at speeds well in excess of electronic components, ultimately being limited by the optical cycle time (*ca.* 1–10 fs). Nonlinear interactions are being studied in numerous laboratories in the subpicosecond time frame and seem certain to lead to novel techniques for pulse shaping (Weiner *et al.* 1988), interaction (Doran & Wood 1987) and transmission (Mollenauer *et al.* 1986). This in turn will allow simple pulse sampling, interleaving, deinterleaving and coding. There being no electronic competition at these rates, optics will be obligatory if such operations are required. However, the needs of the next decade or two lie primarily in processing at multigigabit-per-second rates and in times of tens to hundreds of picoseconds. Here optical logic will have to compete directly with electronic logic if it is to find applications.

The most widely discussed optical logic devices are the bistable optical switches (Gibbs 1985), which have been demonstrated to operate, in threshold logic mode, as dual input AND, NAND, OR and NOR gates with binary on/off optical signals. The devices all rely upon some form of positive feedback to achieve the optical equivalent of the flip–flop, the most common being the use of a Fabry–Perot etalon filled with a nonlinear response optical material. The nonlinear material is chosen to exhibit an intensity-dependent refractive index so that the effective optical path length within the etalon or resonant cavity becomes intensity dependent. Increasing the input intensity can then be arranged to switch the device from off resonance to on, and from reflection to transmission. Reducing the optical intensity produces the opposite effect, albeit with switching at lower input power producing a bistable or hysteresis response.

The typical form of the response curves for such a device are shown in figure 1, which also indicates the signal levels required for operation relative to the switching thresholds. Effective power gain can be obtained by biasing the device close to the switch-up threshold, so that a small additional input power produces a relatively large change in output power. However, operation in this mode carries with it exceedingly tough requirements on the power and device stability and reproduceability (Wheatley & Midwinter 1987).

In common with electronic flip–flops, the bistable optical gates exhibit a switching time that shows a very sharp dependence upon the amount to which the switching input signal exceeds the switching threshold, with very small excess signals leading to very slow switching. In the case of the electronic devices, because they have large intrinsic power gains, it is easy to ensure that the switching threshold is exceeded by a substantial margin. However, in the case of the vast majority of optical devices described so far, having no intrinsic power gain means they are

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

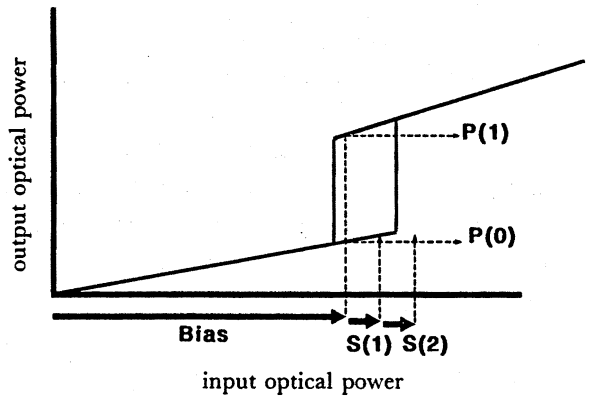

FIGURE 1. Response of a bistable device indicating power settings for use as a dual input AND gate.

constrained to operate within the 'critical slowing down' régime (Wheatley & Midwinter 1986). This presents a further obstacle to their serious use in even simple logic circuits.

Finally, note from figure 2*a*, *b* that, relative to electronic logic, the devices being studied today are generally both slower and more power hungry, with switching energy requirements ranging from a thousand to a million times that of their electronic competition. Such numbers are far too large to be compensated for by other possible advantages, so that there is now little serious expectation that these devices will be used in complex logic circuits in the foreseeable future. However, as we shall see below, there remain good reasons for seeking to invoke optical technology in an intimate manner within an otherwise electronic processor or alternatively, attempting to bypass the 'central routing problem' by dispersing the logical processing out to regions where it can be handled more readily. Hence I shall examine the properties of optical interconnects for use in conjunction with electronic logic while other contributors to this Symposium will address the question of using optics to disperse the coding of data destination to the remote transmitters (using DWDM or CDMA).

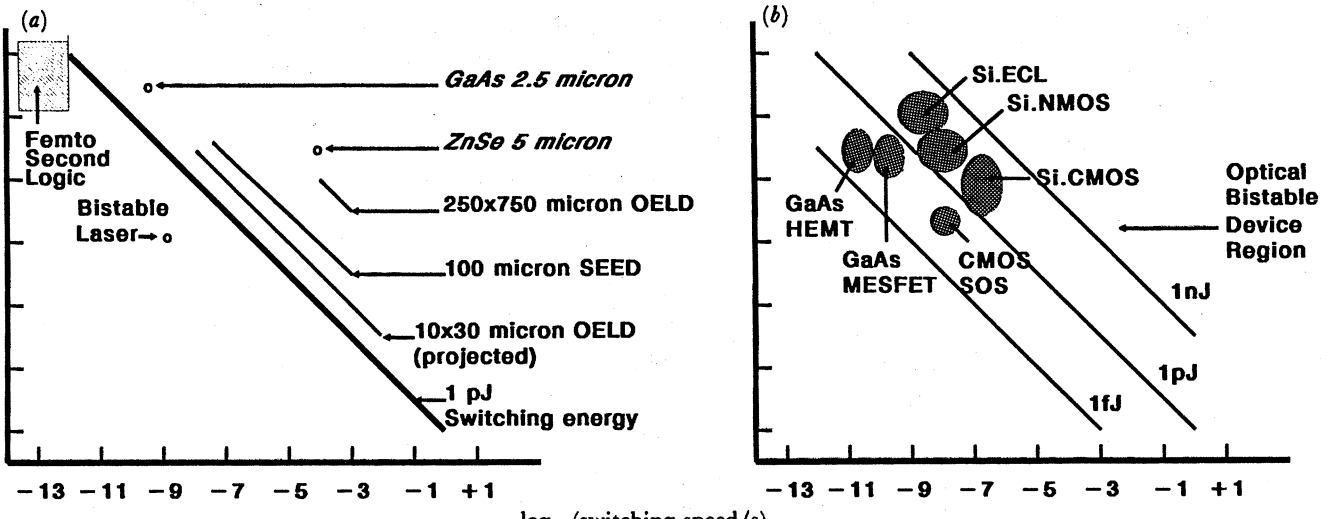

FIGURE 2. Power–speed plot for (*a*) optical logic devices and (*b*) various families of electronic logic devices.

Given that all-optical logic seems to have little to offer to the designer, we turn to hybrid optoelectronic logic and optical interface devices. A growing family of optoelectronic logic devices have been reported (Miller 1987; Wheatley *et al.* 1986) many based upon III–V.MQW material. These rival electronic devices in their sensitivities and may offer similar speeds, but many still suffer from the general stability problem of threshold logic already discussed. Bistable lasers appear to be the major exception because they have intrinsic power gain.

To interface between optical 'wiring' and electronic circuitry, electrical to optical (E–O) and vice versa (O–E) interfaces are required. The former implies a modulated source of light. It is natural to think in terms of a semiconductor LED or laser monolithically or hybrid integrated with the circuit. However, LEDs are very power inefficient and, unless the laser is of a sophisticated design, it is also. Moreover, lasers are critical devices to be grown monolithically along with a wide range of other componentry. Hence, I see great attraction in using a remote light generator to interrogate a modulator placed on the active logic circuit. An ideal candidate is the III–V.PIN.MQW modulator (Miller *et al.* 1985), which can be made both small and fast, and presents itself to the drive circuit as a small capacitance and appears to be well suited to monolithic integration with both III–V and silicon circuitry. For the O–E interface, a variety of fast PIN and Schottky diode detectors exist, as well as the PIN.MQW modulator which can be used as an efficient detector also.

4. Optical wiring or interconnect and its properties

We must consider two forms of optical link that might be associated with our logical circuits, be they optical or electronic, one a planar guide laid over the surface and the other a free space link using imaging components in free space. A planar multimode (step-index) optical waveguide with an index difference of 0.01 would be expected to show a pulse spreading of 0.3 ps cm^{-1} or 3 ps over a 10 cm track. Use of a higher index difference, giving larger dispersion, would be readily countered by some grading of the structure and incomplete excitation of the available mode spectrum, so that we may conclude that bandwidth limitation would be negligible at data rates up to 10 Gb s^{-1} or so. Moreover, this approach will give the shortest delay link between two points on the chip (50 ps cm^{-1} for a refractive index of 1.5) because it follows the shortest physical path at the speed of light. However, by comparison with very large-scale integration (VLSI) metallization, it is likely that optics will be limited, at best, to the equivalent of two-layer circuits because overlaying or intersecting dielectric waveguides is a difficult process to achieve without generating serious cross talk.

Developments such as the silicon 'mother-board' technology (Kawachi *et al.* 1988) seem well placed to provide this type of optical connection, initially between chips precision bonded to the mother-board and perhaps later between different areas of the same chip. A fully satisfactory source and detector technology to work with such guides does not seem to exist although many experimental studies have been reported using existing discrete components in chip-to-chip interconnects.

If we choose to exploit the ability of light to travel in free space (at 33 ps cm^{-1}), then we obtain some advantages but at some expense. For example, if we consider a simple imaging interconnect placed above the chip surface as shown in figure 3, then if we assume that light enters and leaves the devices on the chip normal to the surface, it is axiomatic that the lens must be at least of a diameter equal to the chip diagonal. If we set the chip's linear dimensions as

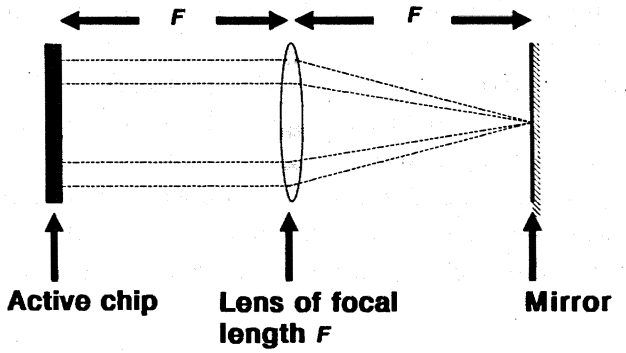

FIGURE 3. Schematic diagram of an imaging zero time skew interconnect placed above an active chip surface to provide 'optical wiring'.

D cm × D cm, then the lens must be at least of diameter 1.414 D. Then, defining the lens to have a fixed f number set by resolution criteria, the focal length is defined as the f number times the lens diameter and the total optical path, chip to mirror and back is four times that if it is to operate as a zero time skew interconnection (see below). Taking the speed of light at 33 ps cm^{-1} in air, an $f2$ optical system and considering a chip with linear dimension of 4 cm, we find that this corresponds to a transit time delay of 1.5 ns and involves an optical system extending a minimum of 22.6 cm above the chip. Since light would travel across the same chip in a glass guide in 200 ps, the 1.5 ns delay implies a very low clock rate if it occurs in the clocked data path. However, with the dimensions of the optical system we have chosen, every path would be identical in delay time, so that a large number of links could be transferred with zero time skew. Note also that a moderate quality optical system of this type could handle in parallel a large number of connections. For example, from a 5.66 cm aperture lens at 11.3 cm distance, the diffraction limited spot diameter in the centre of the field is of order 5 μm (using $\sin\theta = 1.22\ L/D$ where L = wavelength = 1 μm). However, this would also imply severe vignetting of the signals from outer pixels so that a larger value is almost certainly necessary in practice. Note also that this calculation has used the diameter of the central ring of the Airy disc pattern and would thus imply inter pixel interference from outer rings. Accordingly, I suggest a pixel spacing of 30–50 μm is probably appropriate for this example. If extended over a 1 cm square array, it would correspond to as many as 10^5 pixel connections.

Any high-speed digital system must maintain clock and data synchronization between all circuits in a given processor. Thus, the possibility that optics might deliver signals to far and near distant parts of a circuit in time synchronism (zero time skew) is most attractive. If we envisage an object plane and its associated image plane, then for the time delay to be identical between each pair of object or image pixels, it is axiomatic that the phase delay must also be equal. This is easily recognized if we note that in a linear optical system, frequency is fixed and phase delay equals transit time multiplied by frequency ($2\pi x$). Hence, a simple test for a zero time skew optical system is to inject a plane wavefront into it and examine whether the emerging wavefront is also plane. Many imaging systems exhibit this property but, equally, many do not.

The direct imaging of one array to another does not immediately provide a particularly useful wiring pattern, as it corresponds to direct wiring from each element back to it self, apart from a rotation of the array by 180° about an axis normal to its surface. A second identical operation

then brings it back onto itself. It does have the attractive property that it self compensates for linear thermal expansion of the 'chip'.

A more useful wiring pattern in processing is the 'perfect shuffle' (Stone 1971) in which the columns of an array are shuffled, so that annotating the matrix pixels by (i,j), where both i and j run from 1 to N, i maps to i, j maps to k, where $k = 2j-1$ for j in the range 1 to $\frac{1}{2}N$ and to $k = N-2(N-j)$ for j in the range $\frac{1}{2}N+1$ to N. An optical system that achieves this and at the same time achieves nearly zero time skew is shown in figure 4. Other variations have been given elsewhere (Brenner & Huang 1988), but appear to give an unwanted mapping of i mapping

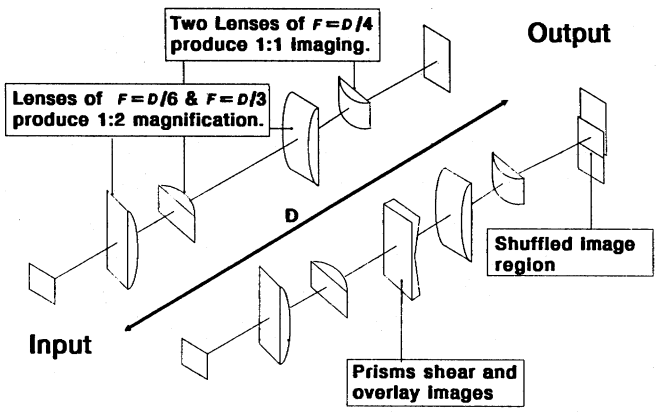

FIGURE 4. Perfect shuffle optical interconnection system.

to $2i$ in addition to the desired j-mapping. Note that the fact that the input data is split into two, overlaid and half effectively thrown away implies a 3 dB insertion loss, which although not absolutely fundamental, requires unattractive additional complication to avoid. Moreover, in many implementations there is a further 3 dB loss arising from the use of a 3 dB splitter to recombine the two signals, giving a total of 6 dB insertion loss.

Another optical wiring pattern known as the 'butterfly' is also widely used in switching matrices. Its optical implementation is shown schematically in figure 5 (Murdocca *et al.* 1988). The input array is optically split to generate three equal power copies, two are shifted laterally (left and right) and the appropriately masked sectors of each are overlaid to form the output array. Once again, power loss is implicit in this implementation because of the power splitting and masking, because even in the ideal case, each output pixel contains at best one third of its incident power (e.g. −4.8 dB). In addition, as a switching matrix based upon butterflies rather than perfect shuffles requires many different butterfly implementations, one attraction of using optics is lost, namely the ability to implement a majority of the wideband connections with a single imaging system. However, it is important to note that a butterfly can always be replaced by a sequence of perfect shuffles (albeit with individual cross points occurring in different spatial locations), so that a matrix using butterflies can be mapped into one using exclusively perfect shuffles but at the expense of some additional wiring steps (see figure 6).

Wiring patterns of this type fit neatly into many parallel pipeline processing operations such a Fourier transformation as well as providing valuable building blocks in switching matrices. However, it is immediately apparent that such a bulk optical system does not neatly integrate with either today's electronic or optoelectronic device technology. The time delay through such

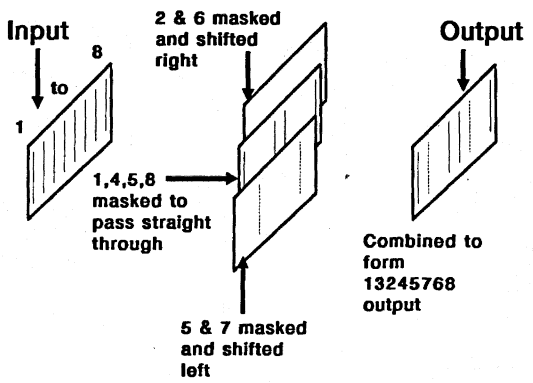

FIGURE 5. Optical butterfly wiring pattern.

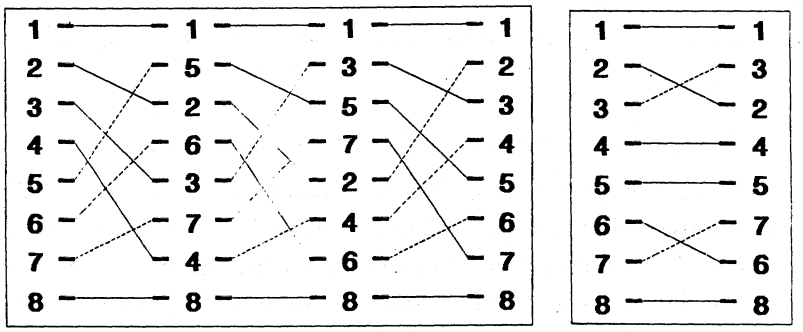

FIGURE 6. Two perfect shuffles reproduce the groupings found in the 8 port butterfly wiring pattern. Two and three shuffles reproduce 16 port butterflies.

a wiring pattern will be substantial and may be an unacceptable penalty to pay for the zero time skew wideband connection it provides. Moreover, if it is desired to concentrate one's active devices on a single chip, then either the wiring pattern must be configured in a loop format, returning the signal to the rear face of the chip and further increasing the time delay or some additional optics must be added, as suggested in figure 7, to allow beam splitting, perhaps using polarization, in front of the chip so that the input and return beams are spatially separated (Midwinter 1987).

Other possibilities exist for imaging interconnects overlaid on top of devices. For example, a holographic Fresnel lens operating in reflection mode could image data in a rather similar manner to that of the lens–mirror combination of figure 3, but it would not provide zero time skew connection. Because the time delay depends upon the physical path length, for two points separated by distance $2d$ on the circuit connected by a lens a distance h above, the path length is simply proportional to $(h^2+d^2)^{\frac{1}{2}}$. Given that $h \gg d$, then this becomes largely independent of d so that again the option exists of trading increased time delay for greater uniformity of delay. With more complex holographic elements, point to multipoint connection patterns could be implemented which could be of great interest in many processor configurations, but the optical system become less clearly distinguishable in its properties to the electronic 'on chip' wiring.

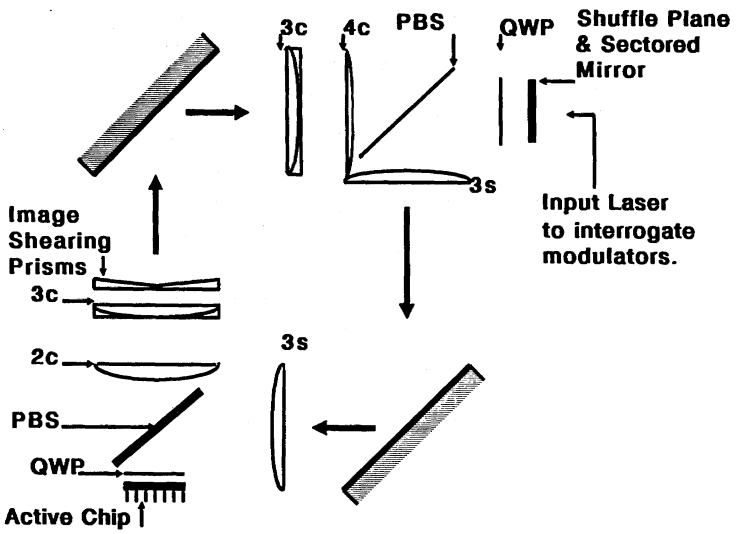

FIGURE 7. Front face wiring system invoking a perfect shuffle and optical power feed.

5. ARCHITECTURES FOR SELF-ROUTING MATRICES

We wish to route very high-speed packet data through a switch. Establishing the path through a complex matrix may require setting anything from one to a large number of cross points. Moreover, as the number is minimized, the complexity of the routing algorithm increases as there are fewer degrees of freedom. In the limit, when there are no free states, changing one pair of inputs can require recalculating and resetting virtually all cross points.

For matrices involving free-space optical interconnections, there appears to be great advantage in selecting a format that is of parallel pipeline form (Midwinter 1987), as shown in figure 8, with the inter-row wiring done by (long time delay) zero time skew optical connection. The total delay through the matrix is then of secondary importance, primary importance being attached to the clock rate at which the data can be clocked through. This in turn focuses emphasis upon the timing precision, which it appears may be one of the few attributes that space optical wiring brings to this problem.

For high-speed packet type switching, a self-routing matrix has great attraction because, in return for a small amount of logic associated with each cross point, each can easily set its own state without needing communication with any other. The destination port addresses must first be transmitted into the input (data) ports. This provides each cross point in the first row with sufficient data to establish its own state. Having done so, the address data then moves via the newly set cross points to the second row where the process is repeated and so on, to the final row. At each cross point, it is only necessary to compare the two addresses entering via its own input ports to establish its setting. Thus the establishment of the path can be done row by row as the addresses are clocked through the matrix, without ever having to externally compute the settings or transmit them to the cross points. An excellent example of an electronic implementation of such a switch is the Starlite switch (Huang & Knauer 1984).

If one injects into the input ports, the complete set of exit port addresses, with no duplication, then a sort algorithm mapped onto the matrix will connect each input in a non-blocking manner to its desired output. This can be done using a Batcher bitonic sorting network

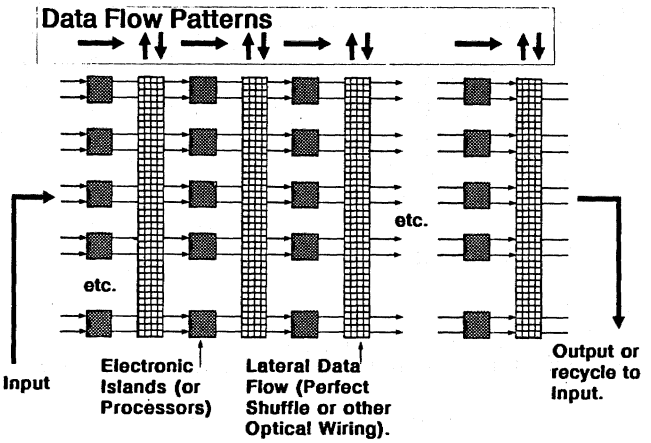

FIGURE 8. Pipeline processor wiring format.

(Batcher 1968), which in turn can be mapped onto a perfect shuffle bitonic sorting network (Stone 1971). However, if some port addresses are absent, then a sort routine simply delivers the data to the output ports in ascending address order, clumping all the 'blank' addresses together starting at the lowest port address. Thus, if signals are directed towards ports 1 and 4 but none to ports 2 and 3, then we can expect to find the port 1 signal at port 1, but the port 4 signal at port 2. Furthermore, if two addresses are duplicated, they are delivered to adjoining exit ports, e.g. two signals for port 1 and one to port 2, then the two packets to port 1 would arrive at ports 1 and 2 whereas the packet to port 2 would arrive at port 3.

Collision problems can be overcome by providing for some data storage and delay, perhaps via fibre delay lines of one packet length that allow the colliding packet to be extracted and fed back into the switch for a second pass. Ensuring packets arrive at their chosen numerical exit port can be achieved either by always transmitting a complete address set with no duplications and generating the addresses that are missing from the input set with dummy packets attached. However, testing the packet addresses before they have been sorted is a time consuming task and it may be preferable to defer this until after the sort stage. Then, it is simpler to detect colliding and missing packets, recode the colliding ones for a second pass and use an expander network to deliver both them and the remaining packets to the correct address.

In both cases, essentially identical self-routing matrices can be used, in the sort case the logic in the cross point has to identify which address is largest, requiring a bit by bit comparison of the binary addresses in most significant bit (MSB) first format. In the expander case, setting is achieved by using successive bits in the address, MSB for the first row to least significant bit (LSB) for the last row (Huang & Knauer 1984). In either case, the logical structure of the cross point takes the form shown in figure 9.

Of particular interest for optical implementation are those implementations that rely upon perfect shuffles for the wiring between every row of the matrix, because a single free-space perfect shuffle optical system can then provide all the matrix wiring outside the individual exchange–bypass modules. This then leads to a concept of the form already shown in figure 7, where the logical cross points are configured in electronic logic on the chip in rows, requiring

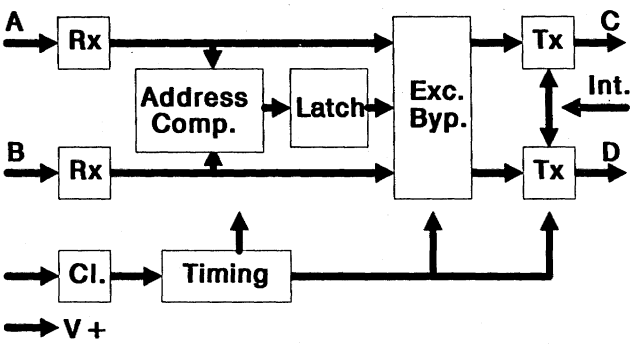

FIGURE 9. Logical layout for the self-routing cross point.

no lateral communication with each other, with all the row-to-row communication done by using a single free-space imaging optical interconnect placed over the chip (Midwinter 1987). To decide whether these concepts should carry any conceivable weight, we need to examine the electronic competition.

6. LIMITATIONS OF ELECTRONIC PROCESSING

One fundamental limitation of electronic chips concerns their ability to communicate with the outside world. This is typically limited to less than 10 Gb s^{-1} with today's pin-out technology, corresponding to a maximum of 200 pins clocked at 100 Mb s^{-1}. Evidently, optical free-space or fibre interconnect technology could have a substantial impact on this figure given a good optoelectronic interface technology (Berthold 1988).

On chip, one's first impression is that electrical communication is rather good. However, a variety of studies (Midwinter 1988; Feldman *et al.* 1988) have independently come to similar conclusions both on power and bandwidth criteria that at clock rates in excess of 0.5–1 Gb s^{-1}, there will be potential advantage in utilizing optical interconnections 'on chip' in place of metal tracks at distances in excess of about 1 mm. The exact break point is sensitively dependent on the electronic and optical technology assumed, the data flow patterns in the circuit and the chosen clock rate.

This has led us to proposals for optically interconnected 'electronic islands' (Midwinter 1988) or 'smart pixels' (Hinton 1988). The island size is dictated by the nature of the data flow required within it. It is defined as the largest electronic circuit which, at the chosen clock rate, can communicate freely internally by electrical means alone, but which benefits from optical communication in communicating with near or far neighbouring islands. Noting that 1 mm corresponds to 1000 μm, and anticipating that the optical interfaces with electronic amplifiers and optoelectronic components might consume 100 μm strips at the entrance and exit to each island, we are left with an electronic processing area of order 800 × 1000 μm per island. By using this space with design rules for 1 μm line width logic and typical cell of 20 × 20 space units (Mead & Carver 1980), we find an upper limit to circuit complexity per island of 2000 gates. Evidently, some of this space would be consumed by electrical connections, say a half, leaving a maximum processor complexity of order 1000 gates, considerably more complex than our basic matrix building block of the 'intelligent cross point' shown in figure 9.

7. Critique

In summary, I have suggested that, in principle, optics can undoubtedly overcome communication bottlenecks in digital electronic systems and moreover, can provide very positive assistance in solving timing the problems in very high clock rate systems. However, most of these gains are purchased at the cost of increased communication delay, which means that only in certain architectures, notably those of pipeline format with no lateral communication between processors in any given row, can the full benefit of such optical techniques be felt. Given such architectures, and switching matrices form perfect examples, very substantial increases in allowable clock rates should be possible. The indications are that the optimum 'electronic island' will be larger than one self-routing cross point, but the detailed design and implementation of such a system has barely been started so that this is a very preliminary conclusion. I am not convinced that 'all-optical logic' has any useful role to play but rather believe that by using optics intimately embedded within digital electronic systems the latter can acquire all the merits claimed for 'all-optical computers'. The sole exception appears to be for ultra-fast, picosecond or subpicosecond simple preprocessing as in a fast multiplexer or demultiplexer.

References

Batcher, K. E. 1968 *Conf. Proc. AFIPS* 1968 *SJCC*, no. 32, no. 2, p. 307. Montvale, New Jersey: AFIPS Press.

Berthold, J. E. 1988 *Conf. Proc., OFC*-88, paper ThI.1. IEEE/OSA.

Brenner, K. H. & Huang, A. 1988 *Appl. Opt.* **27**, 135–137.

Bulley, R. M., Goodman, M. S., Kobtinski, H., Lo, C. N., Loh, K. W., Vecchi, M. P., Brackett, C. A., Curtis, L. & Gimlett, J. L. 1988 *Proc. European Conf. on Opt. Commun.*, ECOC-88, Helsinki, Finland, Sept. 13–17, vol. 1, pp. 345–348. Finland: Professional Electroengineers.

Chung, F. R. K., Salehi, J. A. & Wei, V. K. 1989 *IEEE Trans. Inf. Theory.* (In the press.)

Doran, N. J. & Wood, D. 1987 *Conf. Proc. Topical Meeting on Photonic Switching*, pp. 157–159. IEEE/OSA.

Feldman, M. R., Esener, S. C., Guest, C. C. & Lee, S. H. 1988 *Appl. Opt.* **20**, 1742–1751.

Gibbs, H. M. 1985 *Optical bistability, switching light with light.* Academic.

Hinton, H. S. 1988 *Conf. Proc., OFC*-88, paper ThF.2. IEEE/OSA.

Huang, A. & Knauer, S. 1984 *Conf. Proc. GLOBECOM*'84, IEEE 84CH2064-4, vol. 1, pp. 121–125.

Kawachi, M., Kobayashi, M. & Miyashita, T. 1988 *Proc. European Conf. on Opt. Commun.*, ECOC-88, Helsinki, Finland, Sept. 13–17, vol. 2, pp. 53–78. Finland: Professional Electroengineers.

Mead, C. & Conway, L. 1980 *Introduction to VLSI systems.* New York: Addison Wesley.

Miller, D. A. B., Chemla, D. S., Damen, T. S., Gossard, A. C., Wiegmann, W., Wood, T. H. & Burrus, C. A. 1985 *Phys. Rev.* B **32**, 1043–1060.

Miller, D. A. B. 1987 *Opt. Engng* **26**, 368–372.

Midwinter, J. E. 1987 *Proc. IEE. Pt. J.*, **134**, 261–268.

Midwinter, J. E. 1988 *Phys. Technol.* **19**, 101–108 (part I), 153–157 (part II).

Mollenauer, L. N., Gordon, J. P. & Islam, M. N. 1986 *IEEE J. quant. Electron.* **QE-22**, 157–173.

Murdocca, M. J., Huang, A., Jahns, J. & Streibl, N. 1988 *Appl. Opt.* **20**, 1651–1660.

Stanley, I. W., Hill, G. R. & Smith, D. W. 1987 *IEEE J. Lightwave Technol.* **LT-5**, 439–451.

Stone, H. S. 1971 *IEEE Trans. Comput.* **C-20**, 153.

Weiner, A. M., Heritage, J. P. & Salehi, J. A. 1988 *Opt. Lett.* **13**, 300–302.

Wheatley, P. & Midwinter, J. E. 1986 *Electron. Lett.* **22**, 1134–1135.

Wheatley, P. & Midwinter, J. E. 1987 *Proc. IEE. Pt. J.* **134**, 345–350.

Wheatley, P., Whitehead, M., Bradley, P. J., Parry, G., Midwinter, J. E., Mistry, P., Pate, A. & Roberts, J. S. 1987 *Electron. Lett.* **23**, 92–93.

Discussion

W. J. Stewart (*Plessey Research & Technology (Caswell) Ltd, Towcester, U.K.*). To describe optical interconnects as 'passive' seems negative; they might be better described as 'non-

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

dissipative'. In fact many of the gates in all-electronic chips are concerned with information transfer rather than actual decisions. This is particularly true of high 'fan-out' devices (e.g. neural nets), so that there is less electronics in an optoelectronic chip even though there are no actual optical gates. This is quite fundamental (see, for example, Feynman 1986).

J. E. Midwinter, F.Eng., F.R.S. I agree with the basic premise. An obvious example is the use of a coherent optical system to carry out Fourier transformation in going from the near-field to the far-field radiation patterns. The neural net is also another example where the information is stored in the analogue strengths of a very complex connection pattern. However, it is noticeable that each of these involves analogue processing, albeit perhaps of digital data. A major concern with such optical techniques has traditionally been the limited dynamic range of optical systems because of the relative large photon energy and the consequent granularity of optical signals. However, I agree that there are some interesting opportunities to explore here. I did not do so in my talk because it was specifically linked to 'digital' signals.

Reference

Feynman, R. P. 1984 Quantum mechanical computers. *Foundations Phys.* **16**, 507–531.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES