High–speed digital optical processing in future networks

D. Cotter and J. K. Lucek

Phil. Trans. R. Soc. Lond. A 2000 **358**, 2283-2296
doi: 10.1098/rsta.2000.0647

10.1098/rsta.2000.0647

THE ROYAL SOCIETY

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

High-speed digital optical processing in future networks

By D. Cotter† and J. K. Lucek

BT Advanced Communications Technology Centre, Adastral Park,
Martlesham Heath, Ipswich IP5 3RE, UK

Optical networks have the potential capacity to handle vast quantities of data, but until now the optical domain has been used only as an unintelligent analogue transmission medium. Here the authors advance the concept of 'digital optical processing' in networks to provide a powerful interface between electronic data platforms and the optical domain. Digital optical processing provides the key to increasing the speed and functionality of data networks.

Keywords: optical networks; optical processing; digital optical processing; data networks

1. Introduction

The information revolution has been underpinned by enormous growth of transmission and switching capacity in communications systems. Existing networks consist of digital electronic switches interconnected by optical transmission pipes. Optical point-to-point transmission systems with capacities of hundreds of Gbit s^{-1} are available commercially today, and experimental systems with several Tbit s^{-1} capacity on a single fibre have been demonstrated in several laboratories (Chraplyvy & Tkach (1998), and references cited therein). These high capacities are achieved using wavelength division multiplexing (WDM) to carry many optical channels simultaneously on the same fibre. Single channel bit rates have also steadily increased with time, both commercially and in research laboratories. Transmission speeds of 80–160 Gbit s^{-1} are in a stage of development, and speeds as high as 640 Gbit s^{-1} have been demonstrated (Nakazawa *et al*. 1998; Yamamoto *et al*. 1998). In recent years, there has been great interest (Hill *et al*. 1993; Saleh 1996) in extending the immense capability of optics for point-to-point transmission to build an 'optical network layer'. It was recognized that it is inefficient to convert all data arriving at a node into an electronic format when typically the major part of the data carried on the pipe is merely passing in transit to some other destination. It was envisaged that a mix of traffic with different destinations could be carried on the same optical pipe with each destination allocated a distinct wavelength. In this way, at each network node the traffic intended for the local destination could be separated off from the rest of the transit traffic using passive optical devices (such as wavelength channel filters or wavelength add–drop multiplexers). Thus it would be possible to provide an

---

† Present address: Corning Research Centre, B55 Adastral Park, Martlesham Heath, Ipswich IP5 3RE, UK.

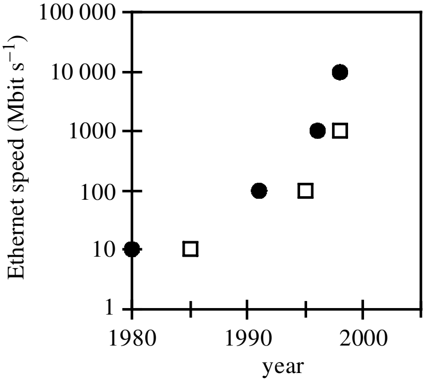

Figure 1. Historical development of Ethernet (●, development work announced; □, first standards ratified).

'express' through-path for the transit traffic, without requiring conversion between the optical and electrical domains (O/E conversion) for processing and routing. The technology needed for this optical routing (the optical wavelength filters and multiplexers, possibly together with some optical cross-connect) would be co-located with the electronic switches. This would result in much fewer O/E conversions and less electronic switching required at each node, with a resultant significant saving in costs. Appropriate setting of the add–drop multiplexers and cross-connects would form end-to-end optical circuits defined by wavelength.

However, within the last two years, it has increasingly been realized that there are flaws with this network model based on wavelength routing. It is an inherently analogue, rather than digital, routing mechanism. This means that there are underlying physical limitations that restrict the cascadability, modularity, and ultimate scalability of this approach. For example, it has been realized that the cascadability of cross-connects is severely limited by accumulative impairments such as cross-talk and noise. This, combined with transmission line impairments (such as group-velocity dispersion, polarization-mode dispersion and the accumulation of optical amplifier noise), means that the wavelength-routed network model has very poor modularity. That is to say, one cannot simply build a network by interconnecting arbitrary numbers of cross-connects over arbitrary distances. When forming the interconnections between network nodes, it is necessary to take account of the origin of the signals and their history—the paths followed, the numbers of amplifiers encountered *en route*, the types of fibre used on the individual links, and so on. (This is in stark contrast with modern digital electronics, which is highly modular, and can be cascaded in an almost arbitrary fashion to produce reliable systems with truly global scalability.) In addition to these physical limitations of optical networks based on wavelength routing, there is restricted scalability from the networking systems perspective simply because a physical resource (wavelength) is being used as the 'address' within the optical layer, so the size of the address space is constrained by the number of wavelengths available. Furthermore, the logical process of address translation requires the cumbersome physical process of wavelength conversion. This is currently difficult to realize in a general-purpose reconfigurable form (any wavelength in to any wavelength out) and can in itself lead to additional impairments such as degradation of the signal-to-noise ratio. As a result of all these limitations, many network developers are retreating from the vision of using wavelength routing in this form in practical, extensive networks. Increasingly the main approach being advocated is to use O/E/O (optical–electrical–optical) transponders at the inputs and/or outputs of the cross-connects in order to regenerate the signals, or to use electronic crossbar switches for routing signals instead of optical cross-connects. It is envisaged that optical cross-connects will play an important role as an extension to the physical layer, but primarily for provisioning, fault protection and managing network paths on a semi-permanent or slowly reconfigurable basis, rather than for dynamic wavelength switching. Thus, unfortunately, the original goal of allowing transit traffic to bypass the electronic domain has not been fully met. Furthermore the great potential for processing information in the optical domain is hindered.

2. The requirements of future networks

Telecommunications networks are currently required to carry rapidly increasing volumes of data traffic in addition to the traditional voice and fax services. This rapid convergence of telecommunications and data networking is forcing a strong shift of emphasis in the telecommunications arena from circuit switching to packet switching. There is an increasingly urgent need for an efficient way of interconnecting data platforms over optical networks. Data platforms such as IP (Internet Protocol) routers are appearing with port speeds reaching those normally associated with core transport networks. Traditional data networking technologies, such as Ethernet, are rapidly being enhanced and developed to handle these higher port speeds (figure 1). It may soon be the case that, for the first time, the port speeds on advanced data platforms will overtake the highest transmission speeds currently available in commercial wide-area transport networks. As an interim step—almost an emergency measure—WDM could be used to provide parallel trunks. That is to say, N physical links (i.e. N wavelength channels) would provide one logical link between router ports. (This would be a curious reversal of the traditional approach, in which several logical links are provided on a single high-capacity physical link.) There are two possible approaches to achieve this.

(i) Each data packet is striped across N wavelengths or N separate fibres (i.e. the first wavelength or fibre carries bits 1, $N+1$, $N+2, \ldots$, the second wavelength or fibre carries bits 2, $N+2$, $2N+2, \ldots$, and so on). However, this would lead to a complex inverse multiplexing problem at the receiving end in order to identify and compensate for the different propagation times of each physical link. In the case where N separate fibres are used, this skew effect would be due to differing fibre lengths. In the case where N separate wavelengths are used, the skew would result from the group-velocity dispersion in the fibre medium and would therefore be a function of distance. Skew compensation would become an exceedingly difficult dynamic problem if path restoration and other management mechanisms are in play in the optical layer;

(ii) Packets from a given flow are always directed along the same physical link, in order to prevent misordering. However, this places extra processing requirements on the router (the router function can no longer be divorced from the physical transmission layer, since the router must keep proper track of which packet flows are associated with each physical link and ensure that outgoing

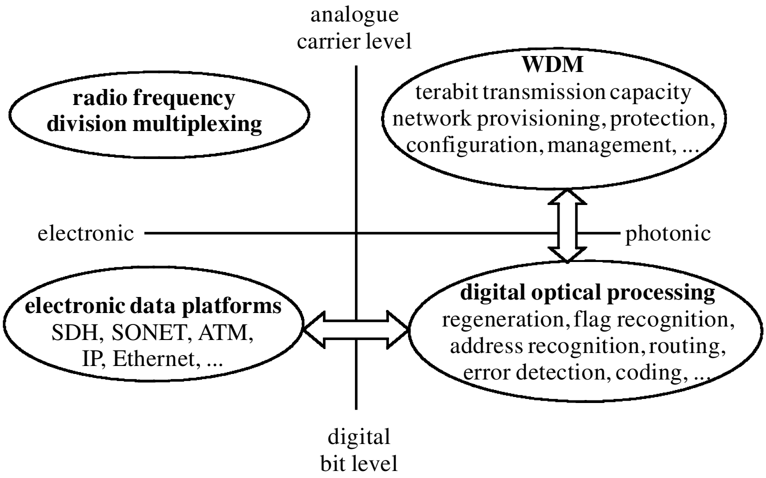

Figure 2. Stages in the development of information handling in communications systems.

packets are directed accordingly). Furthermore, it is well known that a channellized pipe has poorer statistical multiplexing properties than a non-channellized pipe.

If logical port speeds continue to increase at the current rate without a corresponding increase in the speed of the individual physical links, it may soon be the case that the logical ports become excessively channellized. Until only a few years ago, 2 Mbit s^{-1} links were considered quite adequate for data networks. However, few people today would wish to use multiple 2 Mbit s^{-1} links to interconnect advanced data platforms, and it seems very likely that in the future rates such as 2.5 Gbit s^{-1} and 10 Gbit s^{-1} will seem equally restrictive. Rather than attempting to create one logical channel from N physical links, it is highly advantageous to provide a single serial link running at N times the rate. This has consistently been the preferred route in the past; for communication over distances of more than a few metres, it has always proved more convenient to transport data in a serial format. As an example, throughout its development, Ethernet has remained in a single channel serial format (figure 1). (Indeed, the digital optical processing techniques described later in this paper could form the basis of a 100 Gbit s^{-1} (or higher) version of Ethernet.)

From the above discussion, we identify the following requirements for an optical layer:

(i) modularity and cascadability of the network elements;

(ii) ability to allow transit traffic to bypass electronic switches, while avoiding the limitations of wavelength routing discussed above;

(iii) ability to allow high *serial* transmission rates, approaching 1 Tbit s^{-1}; and

(iv) natural interface with packet mode of transport.

3. The evolution of signal processing in communications

Figure 2 summarizes the ways in which the physical signals can be manipulated in a communications system, and puts our discussion into an historical context. The figure is divided into four quadrants. In the days of long-distance transmission on copper cables, it was found that several signals could be sent down the same coaxial cable by using a different radio-frequency carrier for each signal (frequency division multiplexing (FDM); an example of the first, upper-left quadrant in figure 2). This served to increase greatly the traffic-carrying capacity of the cable, but there was no way of processing the signals along the way. This is because the technique is inherently an analogue one, and so the system has no knowledge of the signal format or content. Amplifiers were used to boost the signals and filters were used to separate the carriers from each other after transmission, but there was no way of compensating for accumulated impairments such as noise, cross-talk and nonlinearity.

This was superseded by electronic digital processing (the lower-left quadrant in figure 2). This is far more powerful because individual bits can be directly manipulated. For example, signals can be regenerated so that they can be sent over infinite distances, the elements of the system are modular and cascadable, and switching and routing can be done by performing logic operations on address bits rather than by using physical properties (e.g. carrier frequency). Data from different sources can be easily aggregated into a single bit stream, either on a circuit or packet basis, and separated out at the other end. It is this second quadrant that underlies the whole of the information revolution. FDM came first, historically, because it was easier to implement at the time, but once electronic digital processing became possible it rapidly took over because it is so much more powerful.

The third, upper-right quadrant of figure 2, WDM, is very similar to the first quadrant in that several signals are sent down the same cable (optical fibre in this case) by putting each signal on a different wavelength (frequency) carrier. WDM is currently being deployed widely to increase the transmission capacity of optical fibre on point-to-point links, and is the technology that underpins the explosive growth of communication bandwidth. But as with FDM there is no way of directly processing the signals and, as discussed above, the evolution from WDM point-to-point links to large-scale analogue transmission networks based on wavelength routing presents great difficulties.

The fourth quadrant is *digital optical processing*, a field that is currently in its infancy. As with electronic digital processing, it involves the direct manipulation of bits (rather than carriers), but the difference is that it allows processing functions to be carried out at ultra-high speed in the optical domain, without O/E/O conversion. Digital optical processing allows direct manipulation of data bits ‘on the fly’, as they travel along an optical waveguide at the speed of light. We suggest that the way towards powerful data processing at high speed is to use digital optical processing in a pragmatic combination with digital electronic processing. The processing functions required at the nodes in a data network should be shared between the optical and electronic domains, working in cooperation. There are relatively simple tasks (such as packet header address recognition) that should be performed in the optical domain at the transmission bit rate (up to $1\ \text{Tbit s}^{-1}$), and these would complement more complex processing functions (such as routing) carried out in the electronic domain at lower speed (e.g. at the packet rate; up to around 1 billion packets per second). The tasks performed using ultrafast digital optics would require relatively simple subsystems containing only a small number of elementary optical logic devices, whereas the complex processing at lower speed in the electronic domain could exploit the enormous power and complexity of large-scale integrated electronics. Ultrafast dig-

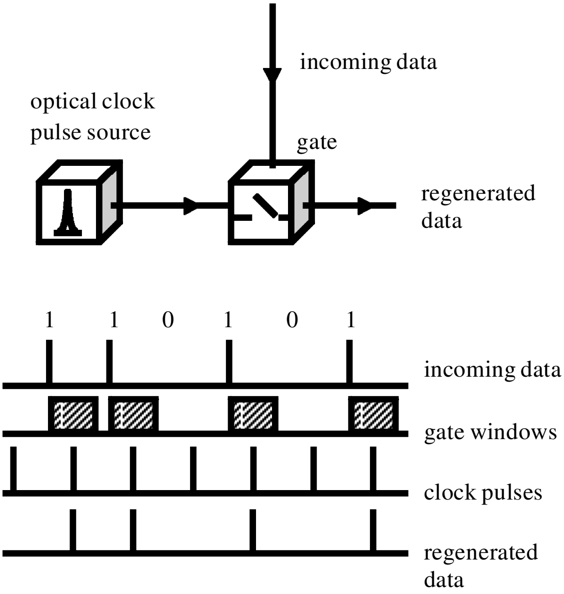

Figure 3. Digital optical regenerator for data represented by return-to-zero pulses.

ital optical processing at present is at the stage of early proof-of-principle work in research laboratories. Demonstrations of a variety of digital optical processing functions have been made recently by combining up to three or four logic gates in various simple optical circuits (albeit so far at relatively low speed, *ca.* 1 Gbit s^{-1}). These have included a shift register with inverter, a regenerative memory with logic level restoration, an optical memory with read and write capability, a binary half-adder, a pseudo-random-number generator, a parity checker, and a full binary adder (Poustie *et al.* (1999*a*, *b*), and references cited therein). Despite this recent exciting progress, it is our view that digital photonic technology is unlikely to compete with the dense integration and rich functionality of electronics, at least for very many years to come. Nevertheless, the pragmatic approach of marrying the strengths of digital optical devices (speed and compatibility with transmission signals) and of electronics (complexity and functionality) is a powerful and realistic prospect, capable of practical implementation within a few years. Moreover, digital optics can provide a highly effective interface between electronic data platforms and the optical domain.

4. Processing in the digital photonic layer

The development of optical networks seems to be following a similar historical path to that of electronics. Currently, optical networks are progressing gradually from the first primitive stage of manually configured patch-panels towards analogue switches (optical cross-connects). The logical next step is the progression from analogue optical switching to digital optical switching and processing. We now illustrate some ways in which *digital optical processing* could be used to meet the requirements of an optical layer listed above.

We envisage the first instance of the use of digital optical processing to be within all-optical in-line regenerators (Lucek & Smith 1993; Cotter & Ellis 1998; Phillips *et* *al*. 1998; Thiele *et al*. 1999). These provide full '3R' regeneration (amplitude restoration, reshaping and retiming) of data signals in the optical domain. By suitably spacing these regenerators along the transmission fibre, in principle they allow infinite transmission distance in the face of impairments such as optical amplifier noise and dispersion. Optical regenerators have been demonstrated in the laboratory, and include an optical pulse source and an optical gate (figure 3). The incoming data bits are used to modulate a continuous train of high-quality pulses from the local source, thus regenerating the original data. Each 'mark' in the incoming data causes the gate to switch to transmission mode for a fixed time (the gate window), allowing a single pulse from the local source to pass through. In the usual design, the regenerator works in a 'synchronous' fashion, i.e. the local source is synchronized to the incoming data stream by some form of clock recovery (not shown in figure 3). In this way the regenerated bits have substantially the same pulse shape, spectral quality, amplitude and timing stability as the local source. Moreover, the regenerator can tolerate a degree of jitter in the arrival time of the data bits, determined by the gate window width.

The digital optical processing field has been transformed by the advent of semiconductor devices capable of operation at 100 Gbit s^{-1} and higher (Cotter *et al*. 1999). Successful all-optical switching using nonlinear semiconductor interferometers with high-speed random-data control signals has been achieved recently. Wavelength conversion at 100 Gbit s^{-1} (Ellis *et al*. 1998*a*) and '3R' regeneration at 40 Gbit s^{-1} (Phillips *et al*. 1998) have been reported. The good cascadability of the devices was demonstrated by incorporating a 40 Gbit s^{-1} '3R' regenerator in a recirculating loop (Thiele *et al*. 1999). It was shown that the switches could be configured to allow 100 Gbit s^{-1} digital optical logic operations (Hall & Rauschenbach 1998). Recently, Kelly *et al*. (1999) obtained error-free operation in a regenerative gate at 80 Gbit s^{-1}.

Looking further into the future, we are working towards networks that process data packets in the optical domain (Cotter *et al*. 1997). Bit manipulation would be carried out 'on the fly' at rates of 100 Gbit s^{-1} and beyond, allowing processing functions such as address recognition, filtering, and error detection to be carried out. This would allow the creation of a *digital photonic layer*, with a possible packet format shown in figure 4. The digital photonic layer header would contain the source and destination layer addresses in addition to any flags or other information fields required within the layer. By way of example, figure 4 also shows a schematic of the first demonstration of binary routing of optical packets based on multi-bit digital address recognition in the optical domain (Cotter *et al*. 1997). In the demonstration, the incoming packet rate was 500 million packets per second and the bit-rate within the packets was 100 Gbit s^{-1}. The address of each incoming packet was examined to see if it matched a particular target address. This would be useful within a node to detect whether a packet is addressed to that node ('for-me/not-for-me' packet address recognition). In effect this binary-decision routing based on digital optical processing would enable the 'express' through-path for transit traffic discussed earlier, but without the drawbacks of analogue (wavelength) routing. In the demonstration the address field of each incoming packet was passed into an ultrafast optical AND gate together with a locally generated bit-pattern corresponding to the logical complement of the target address (i.e. the local digital photonic layer address). The coding scheme was such that if the incoming address differed from the target address, an optical signal appeared at the output of the optical AND gate. Therefore, to effect

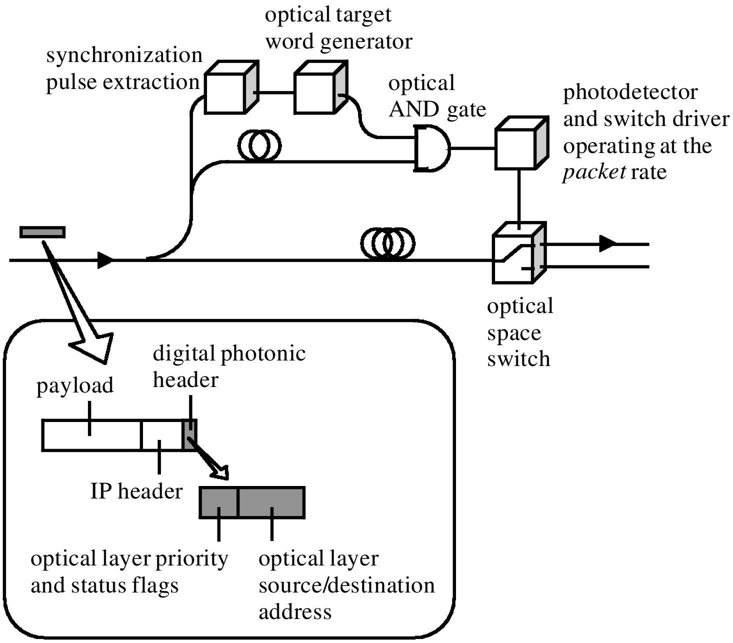

Figure 4. Example of digital optical processing in a digital photonic layer: packet switch based on recognition of a packet header address in the optical domain at 100 Gbit s^{-1}.

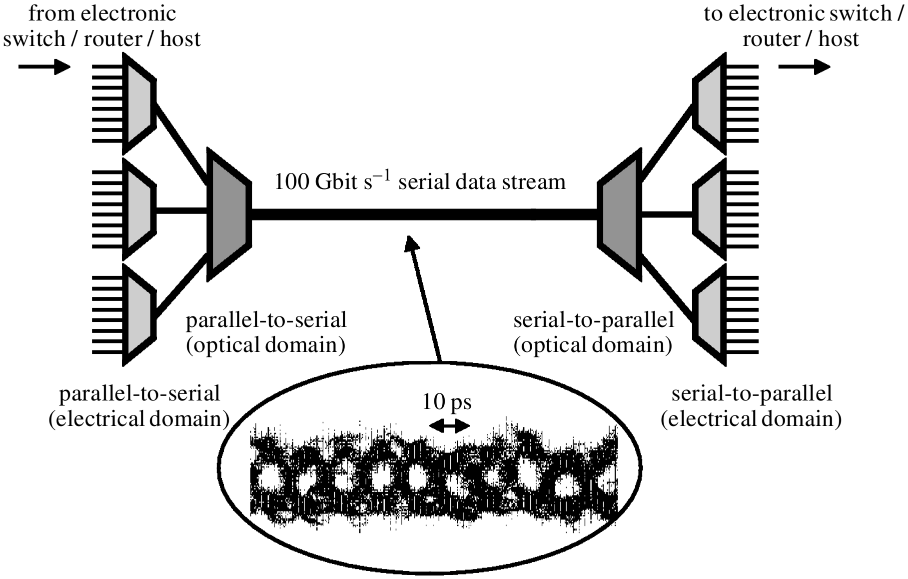

Figure 5. Adaptation between the digital electronic and digital optical domains.

recognition of the target address, it was merely necessary to use a photodetector to detect the output of the AND gate at the *packet* rate (not at the *bit* rate, which is much faster and cannot be resolved by a photodetector). This is a simple example of the principle of using a digital optical device operating at the ultrafast bit rate in combination with electronics that operates at a more modest rate. Moreover, the coding scheme is designed to minimize the amount of ultrafast optical hardware required; only one AND gate is needed, regardless of the length of the address headers. As mentioned earlier, the use of true addresses based on bit sequences within a digital photonic layer, rather than addresses based on physical properties such as wavelength, means that the address space can be very large, unrestricted by the limits of a physical resource (such as the availability of wavelengths).

The technologies to provide the interfaces between the digital optical domain and the digital electronic domain are currently being developed. For example, we recently demonstrated full parallel-to-serial conversion (10×10 Gbit s^{-1} electronic data channels converted to a single serial 100 Gbit s^{-1} optical channel) and the converse serial-to-parallel conversion, using electroabsorption modulator technology (figure 5) (Lucek *et al*. 1998). We can thus foresee the future integration of digital optics with advanced data platforms such as high-capacity routers. As well as address recognition, binary-decision routing and filtering, techniques are currently being developed and demonstrated in the research laboratory to perform other processing functions in the optical domain, such as parity checking and error detection (Poustie *et al*. 1999*b*). There is the future possibility of operating on various fields within digital photonic headers or within the payload, to perform coding or decoding, and to rewrite packet headers for address translation, time-to-live flags, etc.

5. Some impacts on network architecture

Superficially it might appear that all digital processing functions currently performed in the electronic domain would have their direct counterparts in the optical domain. However, there are certain fundamental differences between optical and electronic digital processing. One of the most important differences is that it is difficult to see how to use optics to emulate electronic random access memory. This is because optical signals are travelling waves; therefore it is not possible to capture and hold them in the same way that electronic digital signals (in the form of currents, voltages or charge) can be captured and stored for an indefinite time during which they are continuously available for instant retrieval. Information can be stored in a recirculating optical loop, and it has been demonstrated recently that by introducing a regenerator into the loop the information can circulate for several hours (billions of circulations) without any degradation (Poustie *et al*. (1999*a*), and references cited therein). However, access to the information (via optical taps) is provided periodically in time rather than instantaneously and there is a direct trade-off between the quantity of information that can be stored in a recirculating loop and the speed of access. In the design of a network, these fundamental differences between digital optics and electronics can be overcome, but with significant impacts on network architecture. The lack of fast optical random access memory has an influence mainly on the methods of contention resolution and synchronization, as now described.

(*a*) *Contention resolution*

Contention arises in a packet switch whenever two packets are directed towards the same output port simultaneously. In electronic data platforms this contention is resolved by holding packets in a buffer until the desired output port becomes vacant.

An optical ring network is an example of a network in which the need for optical buffering can be avoided by design. Contention arises when a traffic source wishes to insert packets onto the ring when the ring is already congested. This contention can be resolved by holding the source packets in electronic memory until free capacity becomes available. Once launched onto the ring, an optical packet travels around the ring until it reaches its destination, which is recognized using the binary 'for-me-or-not-for-me' address matching described earlier. In effect, the binary routing can provide the basis of an optical 'front-end' to a conventional electronic router. In addition to ring networks, we are currently developing similar strategies for other networks with regular topology, such as the Manhattan street network (Chevalier *et al.* 1998), which have potential for use in dense high-capacity interconnects, such as gateways or a service provider's 'point of presence'. As in the case of the ring network, all buffering is removed from the optical layer within the network, and instead contention is resolved using electronic buffers at the originating nodes. By using a synchronized switching scheme (called 'clockwork' routing), the routing nodes within the network require only very simple digital optical processing—the 'for-me-or-not-for-me' packet address recognition described earlier—together with digital optical regeneration. Amongst several interesting advantages of this networking approach, the originating node has direct visibility of the network buffer state—there are no relevant buffers hidden from view inside the network—and thus the access and flow control mechanisms can be localized, which is expected to reduce delay and improve handling of bursty traffic.

For optical mesh networks, a method of contention resolution that has received much attention from researchers is to use deflection routing, in which a contending packet is sent deliberately to the wrong output port (Cotter & Tatham 1997). The packet so deflected will then find an alternative route to its eventual destination. In effect this routing strategy uses the network transmission fibre itself as a buffer. This is expected to be quite effective in resolving contention, at least for moderate traffic loads, but has several severe drawbacks: indeterminate and widely varying end-to-end delay, a long delay probability tail, and the need for packet resequencing at the receiver. To avoid these shortcomings, some form of buffering within the switch fabric at the network nodes is needed. The use of optical buffer memories based on switched delay lines is possible, but in practice is likely to be limited to relatively small buffer depths (of the order of 10 packets). Although these small buffer capacities are predicted theoretically to be sufficient to resolve contention in switches handling uniform Poissonian traffic, more realistic traffic models—especially those describing bursty and self-similar traffic—show that buffer depths of several hundreds or thousands of packets may be necessary. Hunter *et al.* (1999) are currently developing ingenious designs for large optically-buffered switches involving multiple cascaded stages of optical switching and buffering to achieve a low packet loss probability. It is envisaged that it will be essential to incorporate digital optical regeneration within these switch fabrics to overcome scaling limitations due to physical impairments such as cross-talk and noise.

(*b*) *Synchronization*

Modern electronic telecommunications platforms (such as the synchronous digital hierarchy, SDH) operate in a synchronous fashion; i.e. the network nodes and sub-systems are synchronized to a single clock which is globally distributed. This can result in the problem that, unless the network is restricted to a simple one-dimensional network topology (such as a bus or ring), data signals arriving at a node from different paths will have different phases. The well-known solution used in electronic switching equipment is to provide a system of buffer storage and pointers to bring signals into phase coherence. However, for networks using digital processing in the optical domain, this solution is impractical because of the lack of suitable memory of sufficient size and access speed. The variations (due to phase fluctuations in the transmitters as well as environmental effects acting on the transmission fibre) cannot be adequately controlled in an uninterrupted fashion. This is due to the restricted physical range of variable optical delay lines, and the limited frequency response of phase-locked loop control systems for source synchronization (because of the fundamental time-of-flight loop delay over extended distances) (Ellis *et al.* 1998*b*). Consequently, there are insufficient degrees of freedom to allow synchronous operation in an unrestricted network topology. To overcome this limitation we have developed an alternative approach—'asynchronous digital optical networking'—in which there is no global synchronization at the bit level (Cotter & Ellis 1998). In this approach, in each network node the clock (used to drive the sub-systems, such as pulse sources, serial-to-parallel converters and switches within the node) operates at a standardized nominal frequency, but is entirely independent and free-running: there is no attempt made to achieve synchronization between the clocks in different nodes. An essential component of the asynchronous digital optical network is a new type of data signal regenerator: the asynchronous digital optical regenerator, which provides '3R' regeneration of optical data packets. This is similar to the signal regenerator shown in figure 3 except that, unlike the synchronous version described earlier, the local pulse source is not forced into synchronism with the incoming data bits. Instead, the asynchronous regenerator works in a reverse fashion, by forcing the regenerated data into synchronism with the local free-running clock. This allows all the incoming traffic at a node to be resynchronized locally. This approach has several advantages, the most important being that it removes the need to maintain synchronization between the different links and routing nodes throughout the network, allowing almost unrestricted complexity and scalability of the network.

(*c*) *Time line*

In summary, figure 6 shows a possible time line for the exploitation of digital optical processing in future networks. The first practical instance of digital optics will probably be in all-optical regenerators for use in high-speed point-to-point transmission systems. This may be followed by the introduction of a digital photonic layer, and the use of binary-decision ('for-me-or-not-for-me') optical packet switching to act as an ultra-high speed front-end to electronic switches or routers. This would be followed by more complex photonic switching nodes (switching between $N \times N$ optical pipes), co-located with the electronic switch, and incorporating asynchronous digital regeneration and processing.

6. Conclusions

This paper has proposed the next radical step in the development of networks to handle a massive increase in data traffic. We have reassessed the requirements for

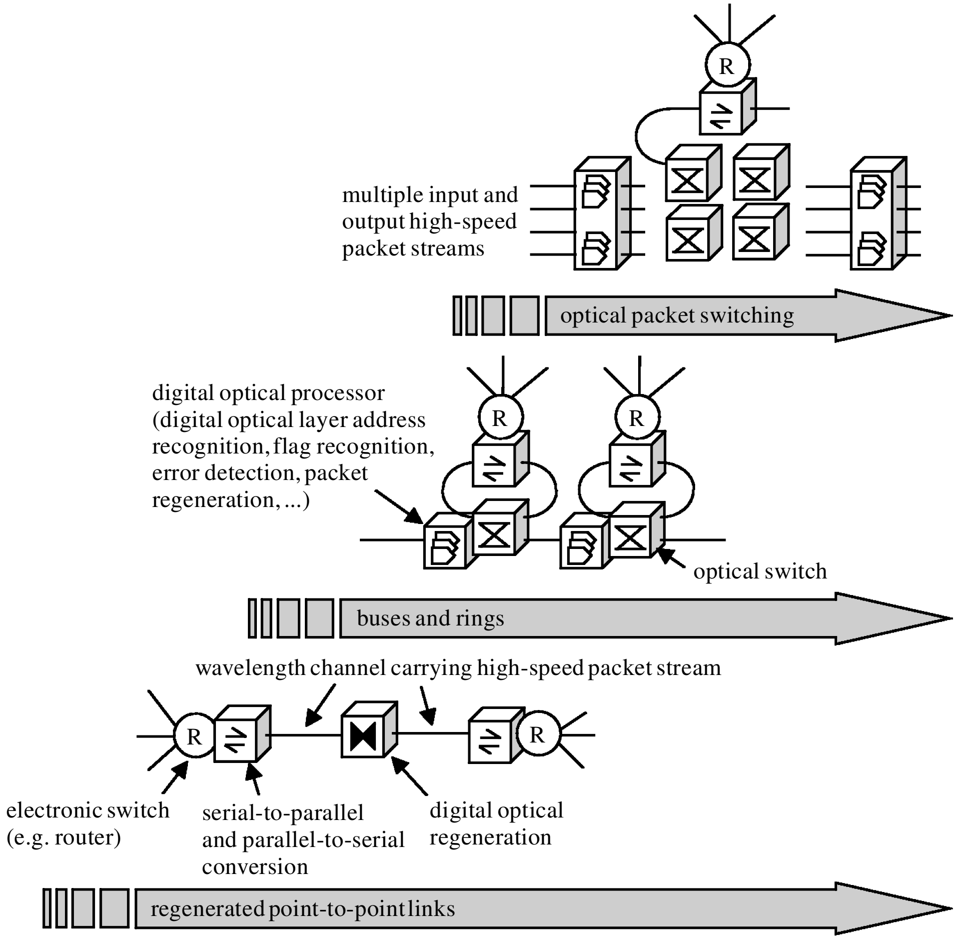

Figure 6. Time line for the introduction of digital optical systems in communications.

an optical network layer from the point of view of the needs and trends of data networking, and have concluded that the existing approaches to photonic networking have significant and fundamental deficiencies. We have introduced the concept of *digital optical processing* in networks (the direct manipulation of bits in the optical domain) to overcome these deficiencies in the future and to provide a powerful interface between electronic data platforms and the optical domain.

There have been several recent advances in developing nonlinear optical techniques for processing serial digital signals at high speed. The field has been transformed by the advent of active semiconductor devices capable of operation at speeds of up to 100 Gbit s^{-1} (Cotter *et al.* 1999). These techniques, working in tandem with electronics, are expected to become important in future high-capacity communications networks, by allowing digital information to be processed 'on the fly' in the optical domain. Devices for optical '3R' regeneration and wavelength conversion are likely to reach commercial deployment within 2–3 years, followed by systems capable of regenerating data in multiple wavelength channels simultaneously. Elementary logic gates will enable recognition of flags, symbols and addresses at the full line rate.

Later, digital optical devices and simple circuits could allow basic packet-level functions to be performed on the fly, such as error detection, header translation, coding and time-to-live mechanisms.

In summary, digital optical processing provides the key to increasing the speed, capacity and functionality of future data networks.

References

Chevalier, F., Cotter, D. & Harle, D. 1998 A new packet routing strategy for ultrafast photonic networks. In *Proc. IEEE Globecom 1998*, vol. 4, pp. 2321–2326. New York: IEEE.

Chraplyvy, A. R. & Tkach, R. W. 1998 Terabit/second transmission experiments. *IEEE J. Quant. Electronics* **34**, 2103–2108.

Cotter, D. & Ellis, A. D. 1998 Asynchronous digital optical regeneration and networks. *J. Lightwave Technol.* **16**, 2068–2080.

Cotter, D. & Tatham, M. C. 1997 'Dead reckoning'—a primitive and efficient self-routing protocol for ultrafast mesh networks. *IEE Proc. Commun.* **144**, 135–142.

Cotter, D., Tatham, M. C., Lucek, J. K., Shabeer, M., Smith, K., Nesset, D., Rogers, D. C. & Gunning, P. 1997 Ultrafast all-optical signal processing for packet switching. In *Photonic networks* (ed. G. Prati), pp. 401–413. Springer.

Cotter, D., Manning, R. J., Blow, K. J., Ellis, A. D., Kelly, A. E., Nesset, D., Phillips, I. D., Poustie, A. J. & Rogers, D. C. 1999 Non-linear optics for high-speed digital information processing. *Science* **286**, 1523–1528.

Ellis, A. D., Kelly, A. E., Nesset, D., Pitcher, D., Moodie, D. G. & Kashyap, R. 1998*a* Error free 100 Gbit/s wavelength conversion using grating assisted cross-gain modulation in a 2 mm long semiconductor amplifier. *Electron. Lett.* **34**, 1958–1959.

Ellis, A. D., Widdowson, T., Phillips, I. D. & Pender, W. A. 1998*b* High speed OTDM networks employing electro-optic modulation. *IEICE Trans. Electron.* **E-81C**, 1301–1308.

Hall, K. L. & Rauschenbach, K. A. 1998 100 Gbit/s bitwise logic. *Opt. Lett.* **23**, 1271–1273.

Hill, G. R. (and 26 others) 1993 A transport network layer based on optical network elements. *J. Lightwave Technol.* **11**, 667–679.

Hunter, D. K., Andonovic, I. & Chia, M. C. 1999 Multi-stage optical buffered switch for IP traffic. In *All-Optical Networking 1999: Architecture, Control and Management Issues* (ed. J. M. Senior & C. Qiao), SPIE Proceedings, vol. 3843, pp. 90–98. Washington, DC: SPIE.

Kelly, A. E., Phillips, I. D., Manning, R. J., Ellis, A. D., Nesset, D., Moodie, D. G. & Kashyap, R. 1999 80 Gbit/s all-optical regenerative wavelength conversion using semiconductor optical amplifier based interferometer. *Electron. Lett.* **35**, 1477–1478.

Lucek, J. K. & Smith, K. 1993 All-optical signal regenerator. *Opt. Lett.* **18**, 1226–1229.

Lucek, J. K., Ellis, A. D., Moodie, D. G., Pitcher, D., Gunning, P. & Cotter, D. 1998 100 Gbit/s parallel-to-serial and serial-to-parallel conversion using electroabsorption modulators. In *Proc. IEEE/LEOS Summer Topical Mtg Broadband Optical Networks and Technologies: An Emerging Reality*, pp. 25–26. Piscataway, NJ: IEEE/LEOS.

Nakazawa, M., Yoshida, E., Yamamoto, T., Yamada, E. & Sahara, A. 1998 TDM single channel 640 Gbit/s transmission experiment over 60 km using 400 fs pulse train and walk-off free, dispersion flattened nonlinear optical loop mirror. *Electron. Lett.* **34**, 907–908.

Phillips, I. D., Ellis, D., Thiele, J., Manning R. J. & Kelly, A. E. 1998 40 Gbit/s all-optical data regeneration and demultiplexing with long pattern lengths using a semiconductor nonlinear interferometer. *Electron. Lett.* **34**, 2340–2342.

Poustie, A. J., Blow, K. J., Manning, R. J. & Kelly, A. E. 1999*a* All-optical pseudo-random number generator. *Opt. Commun.* **159**, 208–214.

Poustie, A. J., Blow, K. J., Kelly, A. E. & Manning, R. J. 1999*b* All-optical parity checker with bit-differential delay. *Opt. Commun.* **162**, 37–43.

Saleh, A. A. M. 1996 Overview of the MONET multiwavelength optical networking program. In *Proc. OFC '96, Conf. on Optical Fiber Communication*, p. 240. Washington, DC: Optical Society of America.

Thiele, H. J., Ellis, A. D. & Phillips, I. D. 1999 Recirculating loop demonstration of 40 Gbit/s all-optical 3R data regeneration using a semiconductor nonlinear interferometer. *Electron. Lett.* **35**, 230–231.

Yamamoto, T., Yoshida, E. & Nakazawa, M. 1998 Ultrafast nonlinear optical loop mirror for demultiplexing 640 Gbit/s TDM signals. *Electron. Lett.* **34**, 1013–1014.