Future high-capacity optical telecommunication networks

Polina Bayvel

Phil. Trans. R. Soc. Lond. A 2000 **358**, 303-329
doi: 10.1098/rsta.2000.0533

Future high-capacity optical telecommunication networks

By Polina Bayvel

Department of Electronic and Electrical Engineering, University College London, Torrington Place, London WC1E 7JE, UK (p.bayvel@ee.ucl.ac.uk)

Since their first introduction in the late 1980s, long-distance optical-fibre transmission systems have advanced from transmission speeds of 140 $\mathrm{Mbit\,s^{-1}}$ for the first optical transatlantic system to exceed 1 $\mathrm{Tbit\,s^{-1}}$ (1 $\mathrm{Tbit\,s^{-1}} = 10^{12}\ \mathrm{bit\,s^{-1}}$) in just 10 years! Laboratory demonstrations of optical transmission at 100 $\mathrm{Gbit\,s^{-1}}$ per channel and use of over 132 channels in a singe fibre are in progress, in the attempt to tap into the huge bandwidth of optical fibres. These developments, in combination with enormous growth in data and voice traffic, have led to many new research ideas for the design of conceptually different high-capacity optical telecommunication networks for the future. This paper reviews some of the dramatic technological developments in architectures, systems and devices, and discusses the ideas of how best to exploit the fibre bandwidth for transmission and processing of information.

Keywords: optical communications; optical fibres; telecommunications; wavelength division multiplexing (WDM)

1. Introduction

Optical communications is a relatively recent research area, but it is advancing at an astonishingly rapid rate (Kaminow & Koch 1997; Lagasse *et al.* 1998). The use of low-loss optical fibres was first demonstrated in the 1970s. The first transatlantic optical communication system, TAT-8, was installed in 1988. It operated at 140 $\mathrm{Mbit\,s^{-1}}$ and used electrical regenerators to amplify, reshape and retime the signals. In more recent years, the capacity of optical networks to transmit data has increased vastly, to tens and hundreds of gigabits and is fast approaching terabit networks.

Of course, this information revolution has important long-run implications for society and the global economy. But these are beyond the subject of this paper. It will be a challenge enough to describe here, in the short space available, the main technical developments in this field, and some of the open research problems.

Technological change is taking place both in the traditional 'optical communications' domain of long-haul, transoceanic systems, and also in landline telecommunication networks around the world. It has provided an impetus for the emergence of a new research area, one that combines many previously distinct research activities. These include fundamental research into optical devices (which includes semiconductor physics and nonlinear optics), research into the transmission of electromagnetic radiation in guiding optical media, transmission systems, multiplexing and coding theory as well as complex traffic and network theory.

Phil. Trans. R. Soc. Lond. A (2000) **358**, 303–329

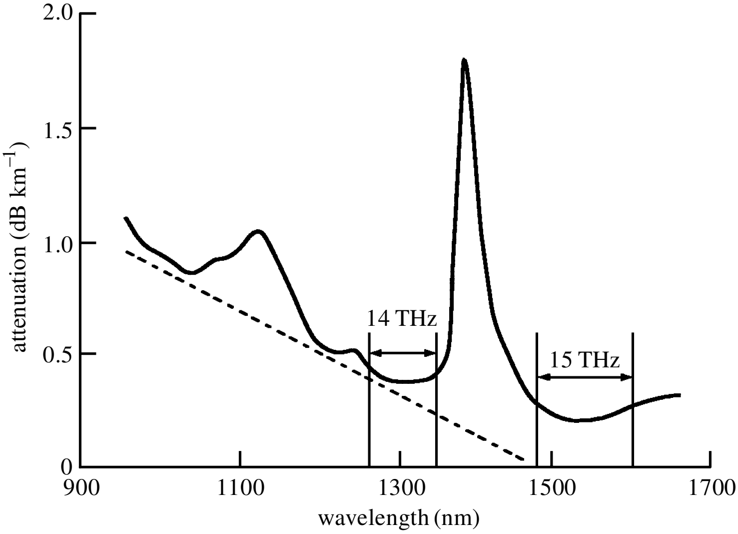

Figure 1. The bandwidth of single-mode optical fibres defined by the regions of low attenuation: the dotted line sketches the fundamental Rayleigh scattering limit arising from random density fluctuations, frozen into fused silica during fibre fabrication; these local fluctuations in the refractive index scatter light in all directions; the peaks in attenuation below 1500 nm correspond to absorption by OH^- ion impurities. Beyond 1600 nm, pure silica becomes highly absorptive (not shown).

The primary limit on the amount of data that can be carried from point to point is the available transmission bandwidth of the medium itself. Single-mode optical fibres are now the ubiquitous transmission medium for the large part of the telecommunications network, except for the last 3 km between the subscriber and the local exchange. The realization that the transmission bandwidth of single-mode optical fibres (figure 1) can provide over 100 nm of optical bandwidth, defined by the spectral region of 1500–1620 nm where optical fibres have very low losses, has stimulated one of the most buoyant research activities. A wavelength range of 100 nm represents over 15 THz of potential frequency space for the transmission data. This does not take into account an additional *ca.* 14 THz (1 THz = 10^{12} Hz) at *ca.* 1250–1360 nm. This spectral region has been temporally abandoned with the development of all-optical erbium-doped fibre amplifiers capable of all-optical amplification at a wavelength of *ca.* 1550 nm. However, the recent research on new Raman amplifiers and new fibre dopants to give amplification of *ca.* 1300 nm is likely to open up this window in the future. Removal of the hydroxyl (OH^{-1}) ion impurities in the silica fibre will reduce the absorption at wavelengths in the range of 1300–1500 nm, to the fundamental limit defined by the Rayleigh scattering, giving a continuous available bandwidth of over 30 THz. What is the best way to use the huge reserve of bandwidth, which exceeds all the radio, high-frequency and satellite spectra by several orders of magnitude?

2. Slicing up the optical bandwidth

Two approaches are clearly possible. The first, wavelength division multiplexing (WDM) allows the spectrum to be sliced up into channels with a different wave-

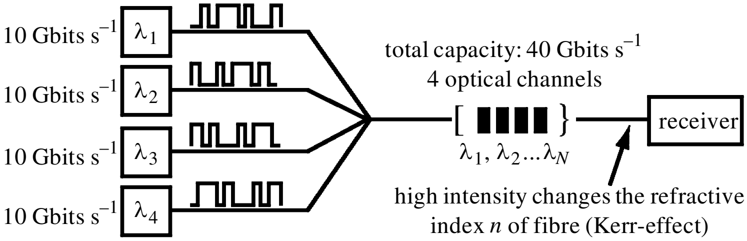

Figure 2. Schematic of a WDM concept, showing how, for example, four different wavelength channels in a single fibre, individually modulated at 10 Gbit s^{-1} each, can give the total capacity of 40 Gbit s^{-1}.

length allocated to each channel (figure 2), carrying a convenient bit-rate, in an analogue manner, similar to the frequency division multiplexing in radio technology (Kaminow & Koch 1997; Lagasse *et al.* 1998; Kaiser 1999). The spectral efficiency with which this bandwidth can be accessed, measured in bits per second per hertz, is becoming important, since it defines how many channels can be 'slotted' within the available spectrum. While the available fibre bandwidth is huge, the requirements for the number of channels is also increasing, to date exceeding 132 channels per fibre. The techniques for improving the spectral efficiency of the systems at the very high modulation rate by using new modulation formats and coding schemes form the basis of yet another growing research area.

Fundamentally, the spectral efficiency is determined by the availability of stable single-frequency sources and demultiplexers/routers able to select individual wavelength channels out of the entire spectral comb. The closer the channels can be spaced, the denser the WDM system is said to be. The more dense the system, the more worthwhile the investment to develop this initially temperamental technology.

The majority of the signals in optical networks are in digital format, with intensity modulation of the light signifying 'ones' and 'zeros': in a straightforward way of either turning a laser on–off or using an electro-optic modulator. The latter acts as a switch transmitting light under a '1' electrical signal and blocking transmission for a '0'. The speed at which light can be modulated is determined by the performance of electronic devices (how fast the stream of '1's and '0's can be generated) and transmission bandwidth of optical components (lasers, modulators and receivers). It is governed by how fast and accurately they can 'translate' an electrical signal into an optical one and back into the electrical domain. Some experiments demonstrating electrical multiplexing capability at 40 Gbit s^{-1} using high-speed InGaAs–GaAs photoreceivers have been reported recently (Takahata *et al.* 1998; Hurm 1998; Yonenaga *et al.* 1998), although the bit-rate that can be conveniently converted from the electrical to the optical domain is currently still 10 Gbit s^{-1}.

Undoubtedly, the modulation rate per channel will increase as research into faster lasers and receivers leads to new devices. However, future increases in the electronic modulation speed will be at the expense of ever tightening tolerances on the residual optical-fibre dispersion to control pulse spreading and ever more careful optimization of the optical transmission medium itself.

The 'granularity', i.e. the convenient electrical modulation rate 2.5–10 Gbit s^{-1} or perhaps even as high as 40 Gbit s^{-1}, deemed high enough to be assigned an individual wavelength is the subject of some debate.

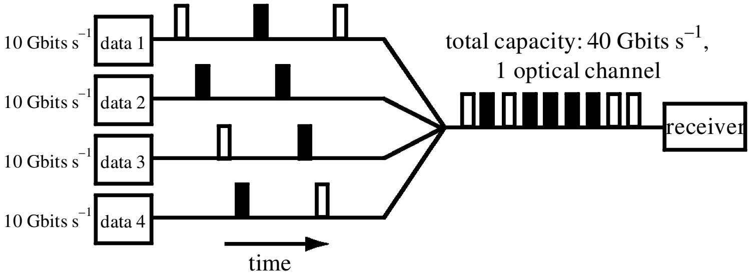

Figure 3. OTDM approach to achieve 40 Gbit s^{-1} in the fibre using a single wavelength. The total amount of data is contained in one optimal channel, transmitted serially over the network and interleaved in time.

A second, competing, network approach, known as optical time division multiplexing (OTDM), would involve all-optical interleaving and regeneration of many lower bit-rate channels, each, say at 10–40 Gbit s^{-1}, and all at the same wavelength, optically multiplexed to 100 Gbit s^{-1} and over (figure 3). The interleaving would be performed by first generating different streams of very narrow pulses, say for a bit-rate of 10 Gbit s^{-1} with a bit-period of 100 ps; the pulse duration would be of the order of several picoseconds. Each stream to be multiplexed would then be separately modulated and delayed, using precise optical delays in free space, optical-fibre or waveguides (2 mm for a 10 ps delay in fibre or silica waveguides) with reference to the initial bit stream, and then time interleaved into a single fibre, generating an $n \times 10$ Gbit s^{-1} aggregate bit-rate.

Paradoxically, both transmission technologies, which increase the amount of data transported from point to point, exacerbate the 'electronic bottleneck'. This eloquently describes the limitation of the electronic switches and routers to process the very high bit-rate signals terminated on their ports. The telecommunication network consists of high-capacity links, connected by electronic processing nodes, switches or cross-connects. As the number of high-capacity links increases, the size of the electronic routing switch must increase too. There are fundamental limitations to this, including electrical 'wire' delays in integrated circuits and backplane electronic interconnects between individual routers within a large switch or cross-connect node. In large routers or switches, these interconnects can be quite long, and, thus, are unable to sustain high electrical speeds. Research into using optics for interconnection within routers/cross-connects has, to date, been unable to yield significant advances, and an electronic switch or cross-connect can, therefore, still only process much lower capacity signals, compared with the terminated rate. For example, a 10 Gbit s^{-1} signal terminated on an electronic cross-connect requires electrical demultiplexing to 64 ports for typical routing rates of 155 Mbit s^{-1}. The growth in electronic ports is also limited by the processing delays associated with computational complexity of route look-up for each 155 Mbit s^{-1} signal, and repeated for all the 10 Gbit s^{-1} channels. Much work has been done in improving electronic router design by speeding up route look-up algorithms. Yet most of these scale poorly with routing table size (Keshav & Sharma 1998; Kumar *et al.* 1998).

So, the potential of using the wavelength domain not just for transmission but for routing proved irresistibly appealing, spurring the work in a large number of research laboratories. In this conceptually simple approach, termed wavelength routing, wavelengths can be used to denote routing destinations, directing high-capacity all-optical signals, transparently, between source and destination nodes without the need for electronic processing. The routing and wavelength allocation would then be predetermined. A number of questions immediately arise, as follows.

(1) How many wavelengths would be required to provide the required connectivity without intermediate optoelectronic conversion or switching? The answer to this question defines the density of the channel spacing, the required stability of the laser sources, and the accuracy of the wavelength selective components: lasers, demultiplexers, filters, cross-connects.

(2) Would some form of all-optical wavelength conversion be required to avoid wavelength contention?

(3) How many wavelengths (and over what distances?) can successfully be transmitted without incurring errors through channel interactions caused by a combination of fibre nonlinearities in long fibre lengths (modest optical power used in transmission of several milliwatts per channel in single-mode optical fibres with core areas of 80 μm^2 give rise to intensities of the order of 100 $MW\,m^{-2}$, at which many optical nonlinearities become significant), chromatic dispersion of the wavelength comb, and the accumulated cross-talk from imperfect selectivity of filters?

(4) How could an all-optical network be managed? After all, only four parameters would be available for monitoring—namely, power, signal-to-noise ratio, optical spectrum and bit error rate—instead of a panoply of electronic monitoring and error correction measures.

Most of the potential limitations are associated with the fundamentally analogue nature of WDM, and the proponents of OTDM were quick to point out that their approach naturally lends itself to digital processing techniques, including digital signal regeneration and bit-serial processing. Both short pulse generation (using, for example, an external cavity mode-locked laser and supercontinuum sources) and regeneration and/or buffering using nonlinear effects in fibres or semiconductors are currently the subject of much intensive research. Still at a rather early stage of research, this approach has very recently seen something of a surge. It offers the potential for a greatly simplified optical amplifier and dispersion management due to single-wavelength transmission. The extension of this transmission technique to network-wide applications also has a number of associated research issues:

(i) devices for generation of stable, wavelength-tunable ultra-narrow optical pulses;

(ii) transmission of short pulses over even medium-haul distances (100–1000 km) requires careful optimization of fibre dispersion, source wavelength and power, as well as polarization and channel delays for correct sub-rate interleaving;

(iii) in the context of extending these techniques to optical-packet networks, how can an essentially non-optical functionality of buffering or memory be implemented optically, given the absence of optical random access memory?; and

(iv) synchronization of packets at nodes.

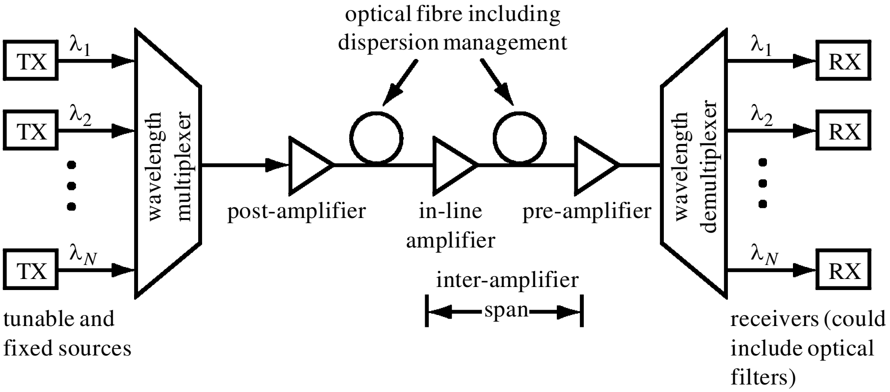

Figure 4. Key building blocks of a WDM transmission system using all-optical amplifiers, and N different wavelength channels (TX denotes transmitter; RX denotes receiver).

3. Evolution of WDM ideas: from concept to reality

As mentioned before, the concepts of WDM were first proposed by analogy with an 'optical ether', or frequency division multiplexing, over 15 years ago (Hill 1988), the ideas explored by this approach languished for many years, until the early 1990s. The enormous growth in demand for high-capacity links over substantial distances of several thousand kilometres resulted in a rapid exhaustion of the installed optical-fibre networks. Since optical-fibre cable installation is so time consuming and expensive, alternative approaches to expand the installed capacity became attractive. The early ideas were 'rediscovered', with the impetus provided by the development of more precise optical components—more stable lasers, all-optical erbium-doped fibre amplifiers—which could now replace electrical regenerators to amplify the entire spectral comb with a single, optically pumped amplifier, and with the emergence from research laboratories of the first practical, compact and stable optical filters that could operate as demultiplexers. The key elements of a WDM transmission system are shown in figure 4.

Although the research work initially separated into the two distinct and mutually ambivalent research fields of 'linear' WDM networks and 'nonlinear' soliton systems, these have now begun to converge.

In a linear or non-return-to-zero (NRZ) modulation format (figure 5), the optical intensity is turned on–off with the bit-period equal to the inverse of the electrical modulation rate or bit-rate. Any appreciable fibre dispersion (due to the wavelength-dependent mode shape and variation of the refractive index) leads to pulse spreading and energy spill over into the following bit-period, resulting in transmission penalties. Dispersion-shifted fibres with a near-zero dispersion are optimal for single-channel transmission to minimize the linear-dispersion-induced pulse broadening. A second distorting effect is the nonlinear phase shift built up on transmission of a high-intensity channel as a result of refractive index change with optical intensity due to the Kerr nonlinearity in the fibre, known as self-phase modulation, and this can also be minimized in low-dispersion fibres. However, the use of low-dispersion fibres is highly detrimental in multiwavelength transmission, where the combination of high optical intensities and low dispersion leads to phase matching between the different wavelength channels, in turn leading to generation of new 'unwanted' frequency com-

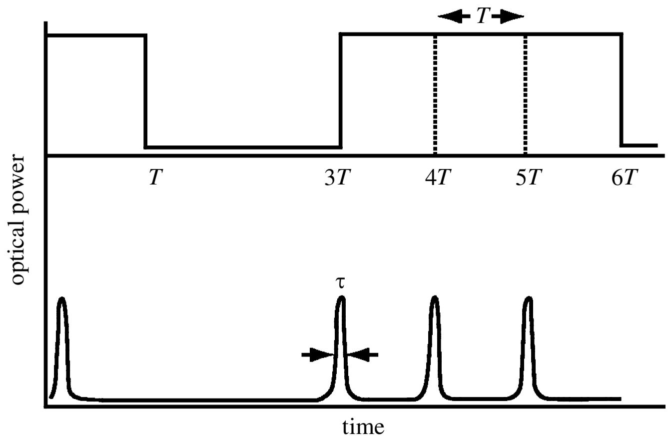

Figure 5. Two different optical modulation formats, NRZ (top) and RZ (bottom). T is the bit-period, equal to 1/bit-rate and the ratio of τ/T could be 0.2–0.5, depending on the system. The RZ format requires sources capable of generating very short optical pulses.

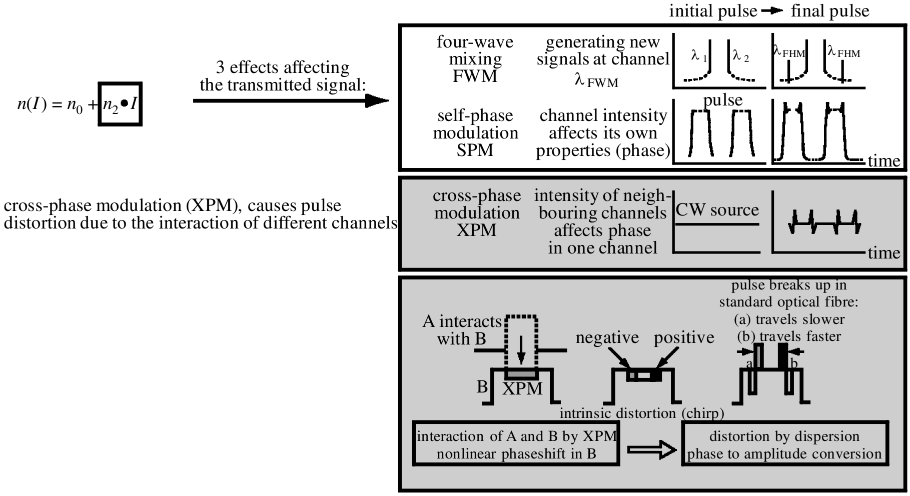

Figure 6. High optical intensities, inevitable in single-mode optical fibres with an 8 μm core radius, give rise to intensity-dependent nonlinear effects. A key nonlinear effect in optical-fibre systems using multiple-wavelength channels is XPM. In the simplest case, the nonlinear phase induced by a high intensity of a modulated channel is converted to variations in amplitude of a CW channel. In a real system, with many channels, all modulated, the amplitude variations become difficult to predict. Two channels, A and B, are considered. A is the interfering channel (pump) and B is the detected channel (probe), which is affected by XPM. If the signal in B is constant, XPM can be separated from SPM and FWM (I is the optical density).

ponents (an effect known as 'four-wave mixing', or FWM). Even more detrimental, the phase modulation of one channel by its neighbours, which can be converted into intensity modulation in the course of transmission, is a result of another unwanted nonlinearity, cross-phase modulation (XPM) (see figure 6 for an illustration of fibre nonlinearities). The investigation of the relative magnitudes and importance of these in different system configurations with different numbers of channels and channel spacing is the subject of much intensive research in many laboratories around the world. Economics seeks an increase in the distance between amplifiers. This, in turn, implies higher power per channel required to maintain the signal-to-noise ratio for each channel to ensure a low bit-error rate. High channel powers inevitably lead to transmission impairments due to fibre nonlinearities, the interplay between which must first be understood and then minimized through the use of appropriate tailoring of the dispersion in the transmission fibre, amplifier design and choice of channel spacing. Use of higher bit-rate per channel with many multiwavelength channels places ever more stringent tolerances on system design.

The competing technology, the transmission of return-to-zero (RZ) pulses (figure 5), whose duration is much less than the bit-period, say by a factor of 3–5, was initially considered a solution to nonlinear interactions, and, specifically, self-phase modulation induced distortion. These soliton-like pulses balance the effects of self-phase modulation and fibre dispersion to allow long-distance transmission without pulse broadening, which is sometimes referred to as nonlinearity-supported transmission. RZ pulse generation itself is the subject of some considerable research, and RZ pulses are a prerequisite for OTDM transmission, since a number of pulses must be time interleaved within the bit-period to generate a much higher aggregate signal, as already mentioned. In actual fact, RZ pulses may or may not evolve into soliton pulses, depending on their energies and the total transmission distance. Strictly speaking, a soliton is a pulse that does not spread in time or frequency on transmission, independent of distance.

Both NRZ and RZ are amenable to wavelength multiplexing, with the RZ format obviously more spectrally demanding, given much shorter initial pulses, which require greater channel spacing.

4. WDM transmission experiments

Before going on to discuss processing and routing issues, a brief overview of experimental details and devices used to achieve the latest very impressive results should be made.

It is interesting that most of the WDM experiments over any appreciable distances have been, to date, carried out in recirculating loops. An ingenious test facility, arguably it is these experiments that provided the greatest impetus for WDM to go from an impressive laboratory experiment to a commercial reality.

Recirculating loops were first proposed by BT Laboratories in 1991 (Malyon *et al.* 1991) to demonstrate long-distance transmission using a much shorter (100 km, for example) link of fibre and one or several optical amplifiers. Such a short distance with a modest amount of control electronics made it into a flexible and universal test bed, where, by changing the fibre length or dispersion, amplifiers or sources allowed the investigation of a large number of different transmission scenarios. Schematically shown in figure 7, the loop is typically built of two switches, connecting a length of fibre L through a 50/50 (or 90/10 coupler). The switches are conventionally acousto-optic modulators, since these have low insertion loss together with a high on–off ratio of greater than 45 dB, to prevent the 'leakage' of light. Initially, switch 1 would be on, allowing the loop to fill for the time $L/(nc)$ (c is the speed of light, $3 \times 10^8\ \mathrm{m\,s^{-1}}$, n is the refractive index of silica, approximately 1.5), after which, switch 1 is closed and switch 2 is opened, allowing the signal(s) to recirculate for a given

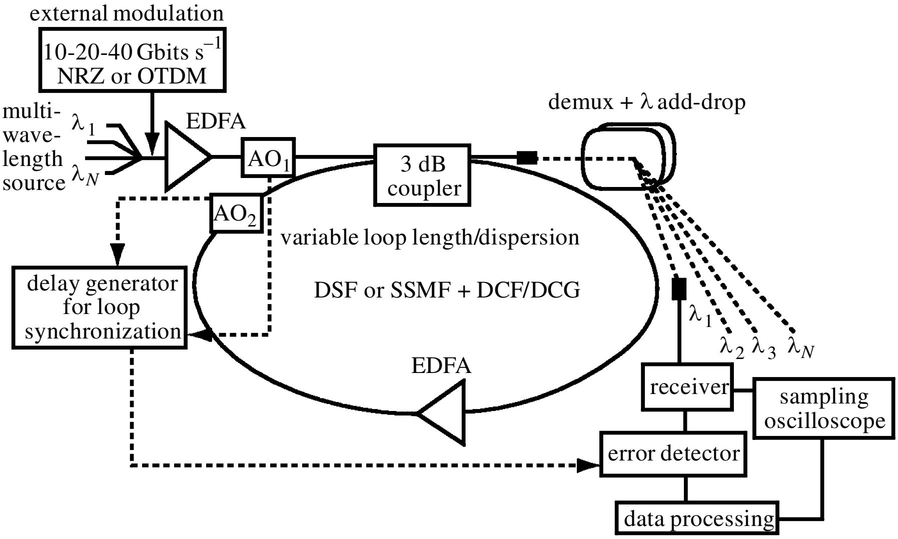

Figure 7. The recirculating optical-fibre loop test bed, widely used as an experimental tool for the laboratory investigation of high-capacity multiwavelength transmission systems and networks (AO denotes acousto-optic switch).

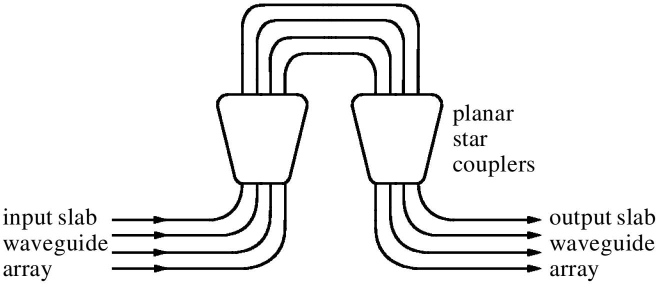

Figure 8. Schematic of an arrayed waveguide grating, which can be used as an $N \times M$ router, or demultiplexer.

number of recirculations N, simulating transmission over $N \times L$ km. The timing delays applied to the switches are controlled by a relatively straightforward electronic delay generator, and the key to the loop stability is the requirement that the product of gain × loss is unity for all wavelengths, preventing any selective amplification of particular channels. An optical isolator must be included in the loop to prevent any counter-circulating, back-reflecting light. The multiwavelength recirculating loop is now an irreplaceable experimental tool (Bergano & Davidson 1995).

A dense WDM network requires stable, precisely defined wavelength sources, either capable of high-speed direct modulation for NRZ transmission or of generating very short pulses for RZ/OTDM transmission. Both must operate stably in wavelength or over a large temperature range. Much research effort has been dedicated to investigating different all-semiconductor, all-fibre and hybrid semiconductor–fibre lasers to fulfil these requirements (Koch 1997).

The development of low-cross-talk wavelength-selective devices has also been key.

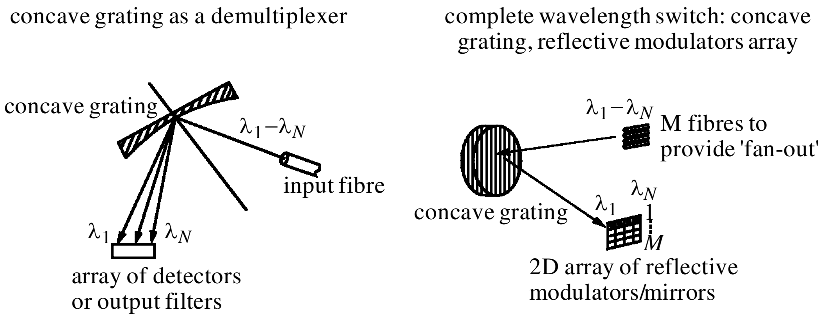

Figure 9. Concave, aberration-corrected diffraction gratings can be used as high-resolution demultiplexers and cross-connects in single-mode optical-fibre WDM systems. Concave gratings provide dispersion and focusing in a single element.

These devices are required to spatially separate or demultiplex the high-capacity multiple-wavelength signals to enable each wavelength channel to be processed: detected, switched or routed to its final destination. This must be done with a low cross-talk: that is the demultiplexed channels must not include spectral components from neighbouring channels and, transparently, to any modulation format. Several approaches have been proposed. One, which captured the imagination, is a planar waveguide device, fabricated on a silica on silicon or InP semiconductor substrate, known variously as an arrayed waveguide grating (AWG), waveguide grating router or a phased array (figure 8). The device is simple yet ingenious (Li & Henry 1997). It consists of an array of waveguides of different lengths with a linear length difference. A multiwavelength light signal will experience a wavelength-dependent tilt of the wave front, resulting in a shift of each wavelength component into a different position, coupling each wavelength into a different output waveguide. With a single multiwavelength input, this device acts as a $1 \times N$ demultiplexer, and with M inputs, as an $M \times N$ router. Although compact and scalable to a large number of channels, the achievable cross-talk is fundamentally limited by the accuracy of the lithographic definition of optical path lengths in different arms of the AWG. Random errors lead to slight inaccuracies in the defined path lengths, resulting in finite and appreciable cross-talk, particularly when the device is used as a cross-connect or router and cross-talk from different input channels adds up coherently, since all the scattering is in the most harmful direction of propagation. It is unlikely that this device can be scaled to a large number of ports for this reason. Moreover, the refractive index of the material used in device fabrication, such as silica or indium phosphide, and, thus, the corresponding spectral positions, depends on temperature. These devices require tight temperature control. The cross-talk and stability performance are particularly important, since, in real network applications, a number of these devices may be cascaded within the source–destination path.

An alternative, if controversial, approach is to use free-space, aberration-corrected gratings, with concave substrates that allow both focusing and dispersion in a single element (figure 9). The long history of diffraction-grating design for spectrographs, spectrometers and telescope applications can be applied to produce devices that have low cross-talk combined with a very high resolution. The aberration control on these devices is somewhat different, since different spectral components must be properly spatially separated, rather than just resolved. Moreover, the aberrations must be controlled tightly to allow focusing of light back into single-mode fibre with a mode diameter of less than 10 µm! However, since these devices are two dimensional, any cross-talk or scattering will be homogeneously distributed within a solid angle, so that the fraction of unwanted light in the dispersion plane of the device will be much reduced. Potentially free-space grating-based devices can be scaled to make much larger routers (Churin & Bayvel 1999).

Although fibre dispersion, that is the variation of refractive index with wavelength, was viewed as potentially limiting, the dispersion limit has largely been solved. The ability to 'tailor' the dispersion map, that is the dispersion of the optical fibre along its length and the ongoing development of new dispersion-compensated, dispersion-flattened and dispersion-slope compensated fibres led to even greater progress (Di Giovanni *et al.* 1997; Gnauck & Jopson 1997). Optimization of appropriate dispersion maps is still a subject of much research, since it has recently emerged that careful balance of nonlinearities could extend both the bit-rate per channel and transmission distances for single-channel and multiwavelength transmission, discussed in the next section.

5. Competition for the highest bit-rate × distance product

Leading research laboratories around the world have competed to report the longest, and ever increasing, transmission distances with the highest bit-rate × distance product using cleverly optimized fibre dispersion, channel spacing and inter-amplifier distance. The first really multiwavelength experiments were reported in 1995 of 20 × 5 Gbit s^{-1} over 6000 km by Bergano *et al.* (1995) and Bergano & Davidson (1995), which in 1998 evolved to 64 × 5 Gbit s^{-1} channels over 7200 km with 0.3 nm channel spacing and 0.14 bit s^{-1} Hz^{-1} spectral efficiency, reaching, in 1999, the highest to date: 64 channels × 10 Gbit s^{-1} over 7200 km spaced by 30 GHz (0.24 nm), spectral efficiency of 0.33 bit s^{-1} Hz^{-1} (Bergano *et al.* 1999). The latter could only be achieved using Reed–Solomon forward error correction (FEC). Longer distance transmission has also been reported by Japanese groups, in experiments using 25 channels × 10 Gbit s^{-1} spaced by 0.7 nm (80 GHz) over 9300 km (Murakami 1998), and in 50 × 10 Gbit s^{-1} transmission over 4000 km with 0.4 nm (50 GHz) channel spacing (Imai *et al.* 1999), and 20 × 10.7 Gbit s^{-1} transmission of 0.8 nm (100 GHz) spaced channels (Taga 1998). At the time of writing, the highest bandwidth 'hero' experiments have exceeded 3 Tbit s^{-1} WDM with the demonstration of 19 × 160 Gbit s^{-1} OTDM channels over 40 km fibre by Kawanishi (1999) at NTT Laboratories in Japan, just beating the previous record of 13 × 80 Gbit s^{-1} OTDM channels spaced by 200 GHz over 89 km (Morita 1998) and 7 × 200 Gbit s^{-1} channels over 50 km, also at NTT Laboratories (Kawanishi *et al.* 1997). The transmission of 132 NRZ channels at 20 Gbit s^{-1} spaced by 33.3 GHz over 120 km (Yano 1996) remains a record in terms of the highest achieved spectral efficiency of 0.6 bit s^{-1} Hz^{-1}! The highest single-channel (by optical time interleaving of 64 × 10 Gbit s^{-1} modulated pulses) transmission of 640 Gbit s^{-1} over 60 km was shown by Nakazawa *et al.* (1998).

Table 1 highlights the conflict between high bit-rate per channel and achievable transmission distances. It is clear that the fight for the longest bit-rate × distance product has been won in long-distance transmission experiments with almost 5 Pbit × km (1 Pbit $= 10^{15}$ bit) this year (1999). The largest bandwidth (that is number of channels × bit-rate per channel) rests with the shorter distance trans-

Table 1. *Bit-rate × distance products for some recent record single- and multi-channel experiments*

author	number of channels	bit-rate (Gbit s^{-1})	distance (km)	bit-rate × distance (Tbit s^{-1} km^{-1})
Bergano *et al.* (1995)	20	5	6000	600
Bergano *et al.* (1999)	64	10.7	7200	4931
Taga (1998)	20	10.7	9000	1990
Murakami (1998)	25	10	9300	2325
Imai *et al.* (1999)	50	10	4000	2000
Yano *et al.* (1996)	132	20	120	317
Kawanishi (1999)	19	160	40	121
Morita (1998)	13	80	89	93
Kawanishi *et al.* (1997)	7	200	50	70
Nakazawa *et al.* (1998)	1	640	60	38

mission of 3 Tbit s^{-1}. All these experiments use some form of nonlinear supported transmission of RZ pulses to offset the effects of nonlinearities and dispersion as well as complex dispersion compensation, including dispersion slope compensation requiring special fibre design. This is likely to remain the subject of much intensive research all over the world for a long time to come.

A single fundamental limit to the ever increasing transmission capacity is difficult to define as it is a function of the system type considered, namely system length, bit-rate per channel, number of wavelengths and linear/nonlinear regime. In the linear regime, where signal powers are low enough not to cause nonlinearities, the limit is determined by two factors. Firstly, fibre dispersion and tolerances on residual dispersion after compensation both in absolute values and in dispersion slope compensation for all the wavelengths. At higher bit-rates, as already mentioned, these tolerances become more and more critical. A second contribution to the fundamental limit is set by the noise accumulation from transmission through many optical amplifiers. Since optical amplifiers simply amplify, but do not reshape or retime, the signal, the noise build limits the total transmission distance. Another factor in this is the finite receiver sensitivity, limited by the shot (quantum) noise and beat noise between the signal and amplified spontaneous emission from the optical amplifiers. These degrade linearly with bandwidth, associated with higher bit-rates, and contribute to limiting the bit-rate × distance limit.

Necessarily low channel powers require short inter-amplifier spans to maintain the required signal-to-noise ratio; this is economically non-viable and pushes the systems into the nonlinear regime. In this case the distortions caused by the combination of self- and cross-phase modulation, and in systems with a very large number of channels, stimulated Raman scattering, also increase with bit-rate, distance and number of channels.

The debate about how best to transmit information—by increasing the number of wavelengths or bit-rate per wavelength channel—rages on. The availability of wide-bandwidth optical amplifiers (Yamada *et al.* 1998), which can provide equal (or flat) gain for all the transmitted channels and can cover the entire single-mode fibre bandwidth of 100 nm, is essential for WDM transmission. However, the future

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

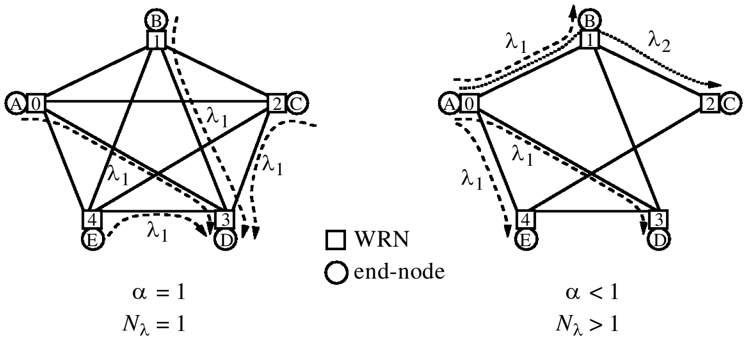

Figure 10. Wavelength-routed optical networks architecture. In this five-node example, ten logical connections must be mapped to a fully connected network; only one wavelength is necessary, while a sparsely connected network needs more wavelengths to avoid contention! How many wavelengths are necessary to interconnect a network with a limited physical connectivity?

of optical communications and networking also hinges on the techniques to be used in processing these huge capacities, with a number of possible solutions.

6. Wavelength-routed optical networks

As already mentioned, the use of wavelength for network-wide routing is an appealing approach. It is an extension of point-to-point WDM, since the number of wavelengths that is required is determined not only by transmission impairments but also by the traffic demand and physical connectivity of the network. A wavelength-routed optical network (WRON) will consist of the interconnection of a large number of point-to-point WDM systems by optical routers or cross-connects, able to route channels according to their wavelength. This has the potential of eliminating the electronic processing at network nodes, one of the limitations on bandwidth growth. A key question, therefore, concerns exactly how many wavelengths would be required to interconnect network nodes on a national or worldwide scale to satisfy a given traffic demand. Research aimed at finding the answer to this seemingly straightforward question has generated a huge number of research publications around the world (Hui & Cheung 1996; Sabry & Midwinter 1994)†.

The question is trickier than first appears, and yet the correct answer determines both the network architecture and the device requirements, such as channel spacing, optical amplifier and router designs and many system and network details. The WRON requires a conceptually different approach to telecommunication network design and operation.

To illustrate the difficulty of the task, consider the problem shown in figure 10, which shows a simple five-node network fully interconnected by bidirectional optical-fibre links with each node connected to four others, so that each node has a unique physical path, disjoint from other node pairs. Each node consists of a transmitter and receiver array and a wavelength-routing node. The transmitters/receivers emit and receive wavelengths corresponding to given lightpaths. Assuming each node needs to communicate to all the others, in a five-node network, then connections need

---

† See also special issues on Optical Networks: *IEEE/OSA J. Lightwave Technol.* **11**, May/June 1993; *IEEE/OSA J. Sel Areas Commun.* **14**, June 1996 and **16**, December 1998.

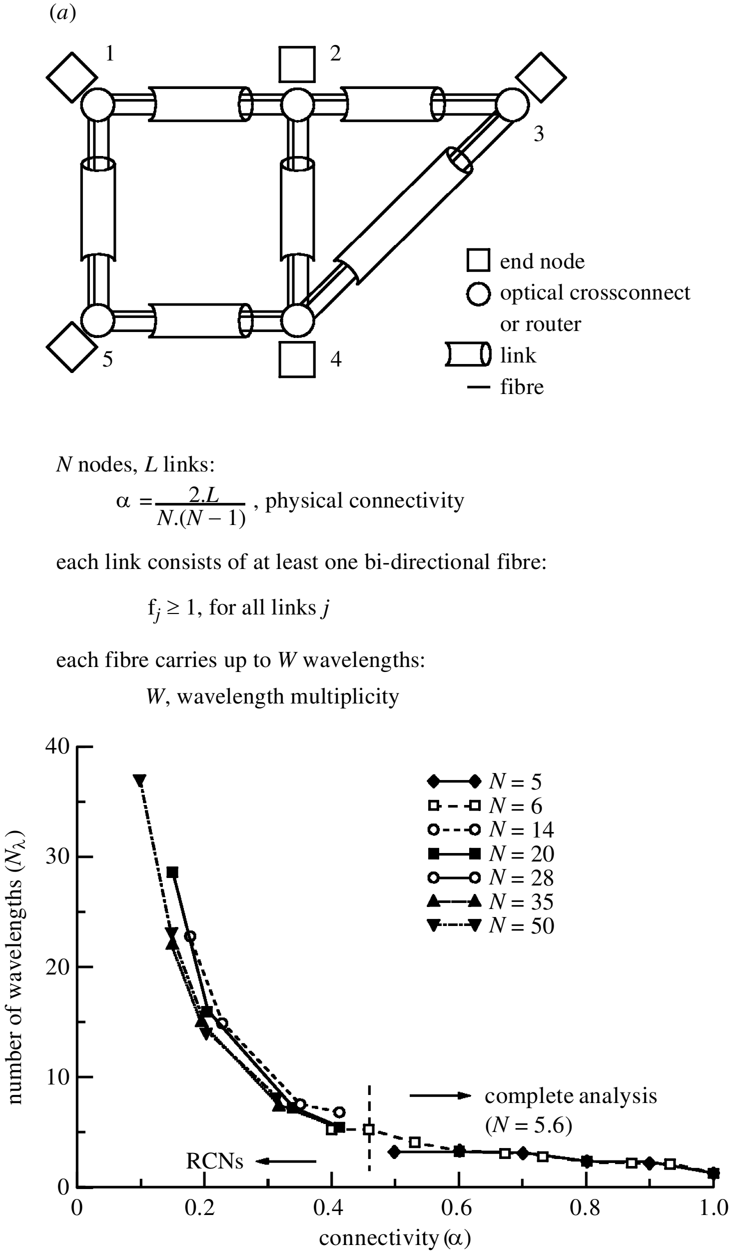

Figure 11. (*a*) Defining main parameters in a wavelength-routed optical network. (*b*) Average number of wavelengths against physical connectivity for networks with a different number of nodes (from Baroni & Bayvel 1997). Note that the number of wavelengths varies with connectivity and is almost independent of the number of nodes! Note also that the same values of connectivity can give rise to many different network topologies. RCN denotes random connected networks: network with arbitrary connectivity. Complete analysis: small networks with five or six nodes can be analysed fully. Note that for the connectivity value of 1, which is a fully connected network, only one wavelength is required (as in figure 10).

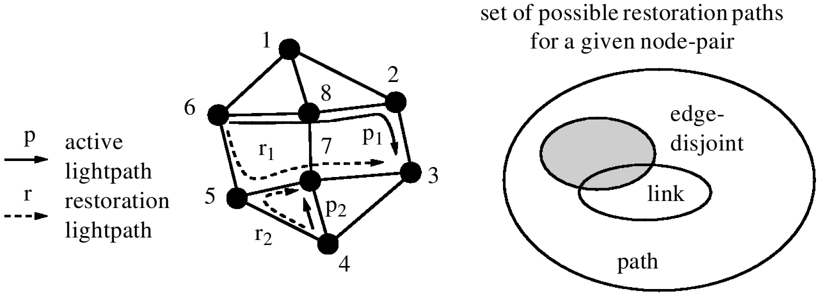

(*a*) edge-disjoint path restoration ($p \cap r = [0]$)

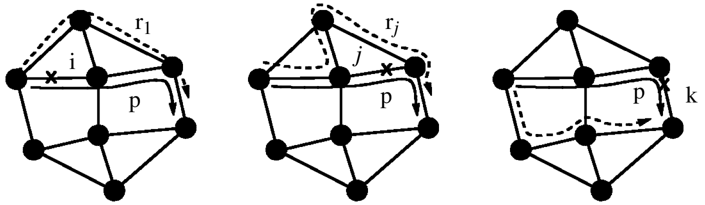

(*b*) path restoration

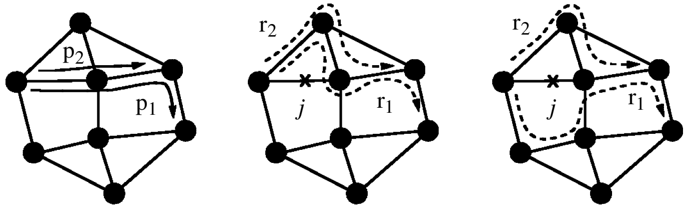

(*c*) link restoration ($p \cap r = [j]$)

Figure 12. Different protection schemes for rerouting wavelength channels in case of a physical link failure (e.g. a cable cut). In each case, a fibre link carries multiple wavelengths and a cable can consist of many fibres. A link failure can be catastrophic, and efficient protection algorithms using a minimum number of extra wavelength channels are vital. Path restoration: the least restoration capacity required (from Baroni *et al*. 1999).

to be made and the same wavelength could then be used for all the connections. However, in a real physical network, the connectivity is far from full, and the required logical connections must be mapped onto available physical connectivity, without different channels with the same wavelengths transmitted simultaneously in the same physical link, referred to as wavelength contention. Moreover, the actual physical connectivity is far from regular and the calculation of the number of wavelengths required if the connectivity were less than 100% is the key issue. The tricky part arises since, given a certain number of links, lots of different physical topologies can be constructed, each requiring a different number of wavelengths. Ideally, one would aim to design a network needing the fewest wavelengths. Figure 11 illustrates the relationship between the number of wavelengths and connectivity. Ultimately, the problem of analysis converts to that of synthesis of good topologies.

An even more difficult question is that of what happens in case of link failures?

Each link would now possibly be carrying a large number of wavelength channels or lightpaths, which would have to be restored by rerouting on different, surviving links. How best should these lightpaths be rerouted to require the least number of extra wavelengths without contention? Should the restoration technique simply bypass the failed link or should a signal be sent to the originating node transmitters to reroute all the lightpaths along pre-calculated alternative paths? Figure 12 shows the possible approaches.

Another group of sub-problems considers the case in which the number of wavelengths per fibre was fixed and the problem relates to minimizing the total number of fibres for a given topology. In both cases, the questions of adding traffic growth to existing lightpaths complicates the analysis.

The problem of wavelength allocation and/or routing preoccupied many groups of researchers in a new field—optical network design—which combined complex techniques borrowed from discrete mathematics, graph theory and combinatorics with the understanding of optical communication principles (Baroni & Bayvel 1997)†. In many cases the problem formulations were too complex to be solved analytically or using integer programming methods, and required formulation of heuristic (or automated 'rule of thumb') algorithms.

Associated with the network design questions were questions of how best to design and control an optical cross-connect or router. Should they be fixed or reconfigurable (see figures 13 and 14) and how best to design them. A battle in the literature and at conferences concerned the necessity of wavelength-conversion devices, which would change the wavelength of a given lightpath along the path, in case wavelength contention occurred. Proponents argued that the use of wavelength converters would make wavelength allocation more flexible, ignoring the fact that one per wavelength per fibre would be required, resulting in a large number of wavelength converter devices requiring monitoring and control. Should these devices be all optical, based on nonlinear effects in semiconductors to generate a new frequency; or optoelectronic, operated simply by detecting the signal electrically and then retransmitting on a different wavelength (somewhat defeating the all-optical routing approach!).

Despite the buzz of activity in this area, the difficulty of achieving consensus has slowed the implementation of these networks. The real difficulty can be traced to the absence of a good reliable optical switch that can be scaled to a large number of ports.

Several techniques are nominally possible, yet none has been shown to give reliable operation for thousand of ports. Research papers have reported 32 × 32 port routers, far short of what is required. So why the difficulty? The story of optical switching is long interspersed with undermined credibility, dating to the days of the hyped-up all-optical computing. Despite the recent dramatic advances in optical technology, optical logic devices, when compared with electronics, are still relatively primitive: they are relatively bulky, poorly integrated, have limited cascadability and are power hungry.

Three recent approaches have shown the most promise. The first is a switch consisting of an array of tiny, perhaps $20 \times 10\ \mu\mathrm{m}^2$, silica micro-machined devices, acting as a two-dimensional array of paddles either transmitting or blocking the beam (Lin *et al.* 1998). Shown in figure 15, these devices are potentially simple and wave-

---

† See also special issues on Optical Networks: *IEEE/OSA J. Lightwave Technol.* **11**, May/June 1993; *IEEE/OSA J. Sel Areas Commun.* **14**, June 1996 and **16**, December 1998.

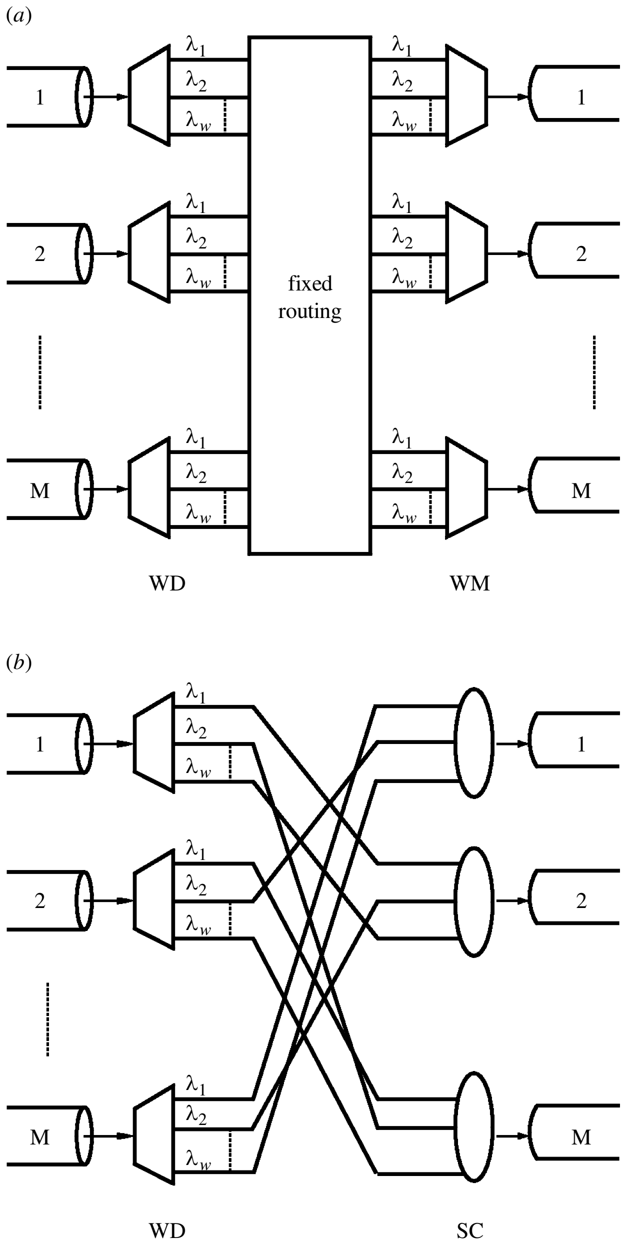

Figure 13. (*a*) Functional blocks of a fixed wavelength router or cross-connect. (*b*) An example of fixed routing. WD is the wavelength demultiplexer; WM the wavelength multiplexer; and SC the star coupler.

length independent with a high contrast ratio—this technology appears fundamentally unscalable beyond tens of ports—since free-space propagating light beamed through the array of paddles must be recollimated periodically, and as the number of passed paddles varies depending on the route, the collimation problem appears insoluble.

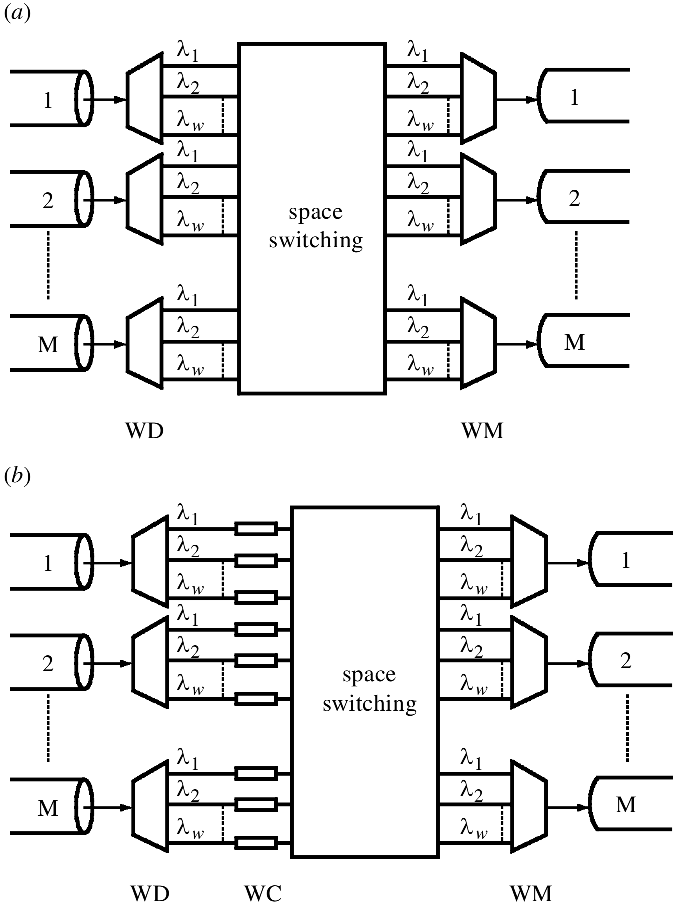

Figure 14. Functional blocks of reconfigurable wavelength routers: (*a*) a router allowing simple reconfiguration of wavelength routes; and (*b*) a router allowing change of wavelength along a particular route. WD is the wavelength demultiplexer; WM the wavelength multiplexer; and WC the wavelength converter.

The second approach relates to cascading of AWGs, described earlier, interspersed by switching elements, such as amplifiers, which would either pass or absorb (and block) the appropriate wavelength. The problem with these has already been alluded to; it relates to the build up of cross-talk with the number of ports. In the context of an analogue WDM network, such cross-talk build up is hard to counteract. Hence, for this approach to be scaled to large networks, some electronic regeneration, judiciously scattered within the network, appears inevitable. The emergence of this idea has led to the term 'islands of opacity in the sea of transparency', with the size and number of the 'islands' the subject of yet more debate.

An alternative, third, proposal takes us back to the free-space device, which would be designed as a switchless $N \times M$ router—with a switching functionality of emitting and controlling wavelengths displaced to the periphery of the network, since the router itself is transparent (Churin & Bayvel 1999)—a possible design is shown in figure 16.

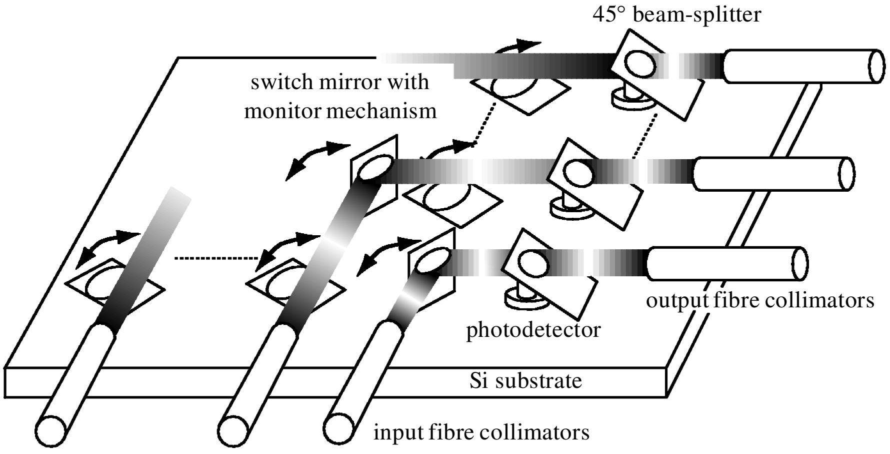

switch mirror: includes operation monitoring mechanism
- monitors configuration of the switch fabric

photodetector: detects signal states and quality, verifies switching path

Figure 15. Example of a free-space silicon micro-machined cross-connect using surface micromachining technology (from Lin *et al.* 1998).

Ultimately, the main advantages of WRONs, their ability to interconnect high-capacity pipes in the top layer of the network, has been questioned by the re-emergence of the packet networks. The new challenge is to understand how packets, rather than continuous high-capacity signals, can be routed across the telecommunications network.

7. Packet networks: the ultimate application of optical networking?

WRONs and high-capacity WDM transmission are the solution for the electronic bottleneck in the transport network. Electrically multiplexed signals are composed of millions of lower-speed voice and data signals before being transported in high-capacity pipes. There are a lot of 'layers' between the user layer and the capacity-rich optical transmission layer. The reason for this is that the high-capacity networks serve millions of users, and per-user cost is relatively modest due to this sharing. However, the cost of providing high bandwidth to the individual user would be costly, and how could the provider ever see the return on the investment? The resources at the user end are likely to be limited for the foreseeable future. One way to help solve the access–transport mismatch is a better utilization of available resources through packet networks. In this way, any transmission is not continuous, thus tying up the entire channel, but is split up into smaller units, or packets, of either variable or fixed length. Each packet has the address of the destination node and a sequence number and is sent into the network to be routed by a number of intermediate packet routers using the addressing information in each packet. At the intermediate or

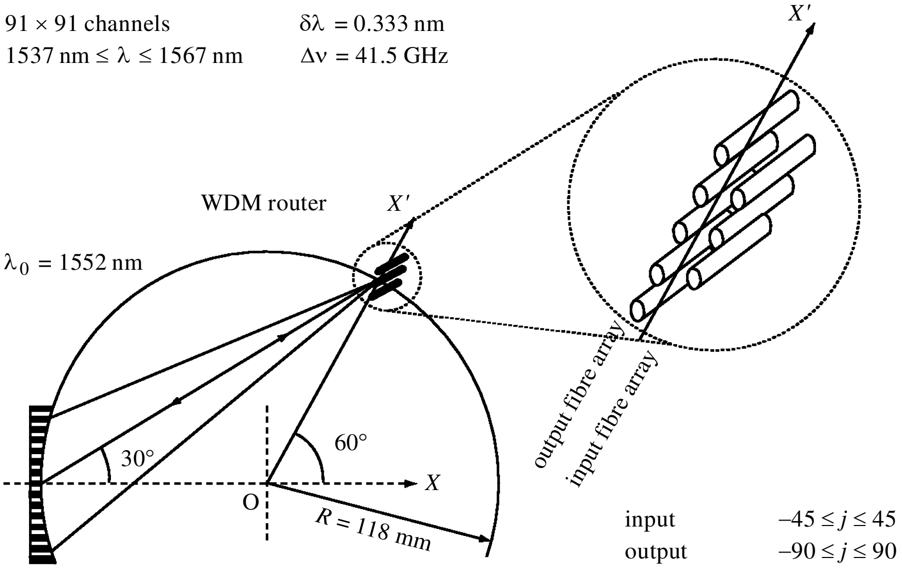

Figure 16. A schematic of an aberration-corrected free-space concave grating switchless router/cross-connect: an example design in which a single concave grating acts as a 91 × 91 wavelength router by dispersing and focusing wavelength channels into individual fibres.

transit routers, there may be packet contention, and delays or buffering is required. At the destination node the original data are reassembled back from sequenced packets. If there are too many packets for the routers, buffers or packets become damaged on transmission through errors or noise, they are said to be lost and may have to be re-sent. Instead of a bit error rate, packet networks have a packet loss rate. The Internet is an example of a packet-switched network.

Originally, packet networks were envisaged only for data transmission, because of the difficulty in ‘packetizing’ voice, as it is rather delay sensitive! However, over the last couple of years, huge progress in packetizing voice signals using compression and speech encoding techniques and improved router design reducing delays through the network have been made and packet-switched networks are seen as a viable (and better) alternative to conventional circuit-switched schemes to improve the use of scarce resources at the user level.

What are packets? They are, in fact, simply bursts of binary data of a predetermined length, say 500–1000 bits, with headers, containing the packet’s destination address, of say up to 10% of the payload (for example 100 bits). Each packet could be carrying a different number of bits at different internal (to the packet) bit-rates. Small packets can be combined into larger packets and routed between users, first in the electronic domain and then in an optical domain. Each channel can then be used more efficiently through better multiplexing or fitting of packets into the available capacity. It has been suggested that, as for the electronic-packet switching, at the user end of the networks, optical-packet networks can help in the ultimate quest towards transparency, providing optical equivalents of electronic-packet networks: namely, buffering, synchronization and routing functions, including header recognition and processing.

In terms of transmission quality, packet-based networks are even more demanding.

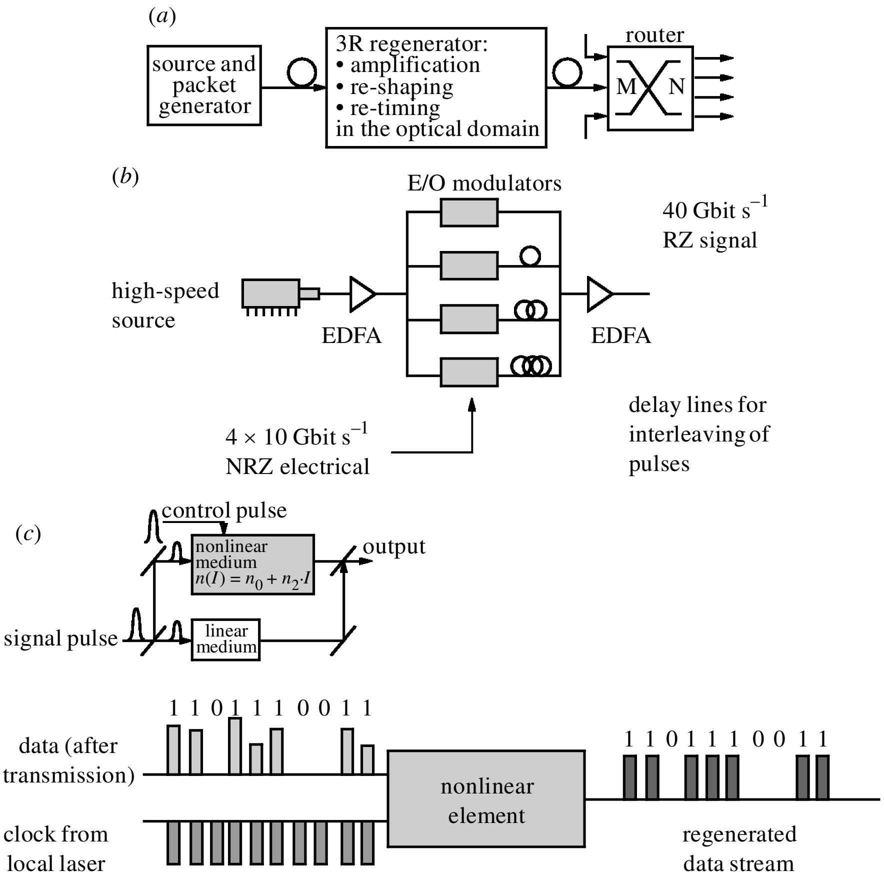

Figure 17. (*a*) Key building functions needed for digital optical-packet networks to become a reality. (*b*) An example source for the generation of packets at, say, 40 Gbit s^{-1} (E/O denotes electro-optic modulators). (*c*) Packet regeneration using a nonlinear element.

A bit error rate of 10^{-12}, which would be sufficient for high-quality long-distance transmission and would indicate a corruption of 1 bit in every 10^{12} bits, may not be sufficient for a packet-loss rate of the same magnitude (10^{12}), as a single corrupted bit could, in fact, corrupt the entire packet! Arguably, without any FEC techniques, packet networks are more sensitive to transmission quality.

In terms of the earlier discussion on WDM/OTDM, packets could either be transmitted on high single or multiwavelength OTDM rates, sequentially *or* in parallel, using different wavelengths. In fact, bit-interleaved TDM systems are isomorphic to WDM. The relative merits of the two approaches have yet to be determined, and the means of generating optical packets from electrical tributaries representing different services and requiring different qualities of service is far from defined. Yet the general area is being seen as an exciting one with much promise (Cotter & Ellis 1998; Cotter *et al.* 1997).

The functionality of an optical-packet network, independent of exact implementation technology, must, nonetheless, include several key components. These functions are shown schematically in figure 17.

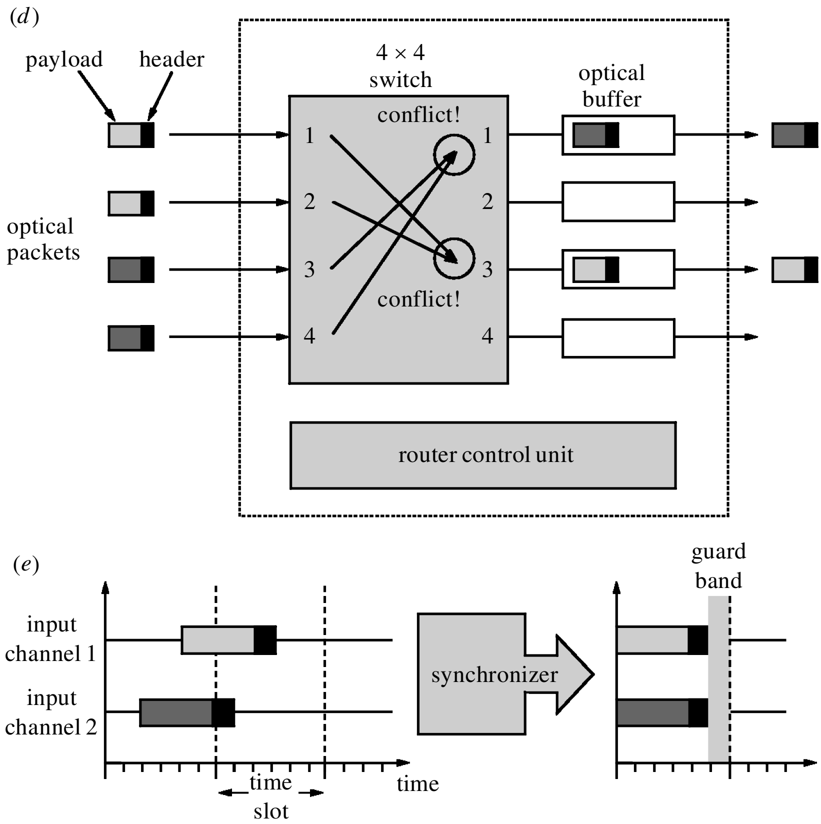

Figure 17. (*Cont.*) (*d*) Optical delay and buffering to avoid contention. Conflict: packets wish to access the same output in the same time slot. Solution: congestion is solved by buffering and delaying packets. (*e*) Synchronization is needed to prevent packets being 'chopped up' at nodes.

The most important is buffering (Hunter *et al.* 1998; Chan *et al.* 1998): the ability to delay packets during header recognition and processing and in the process of switching to avoid packet contention. Since truly random access memory does not exist in optics, the only way to introduce buffering is through precisely set optical delays; typically through a recirculating fibre delay line, where a packet can be delayed by a multiple of the round trip times or, functionally the same, an array of waveguides of different lengths. To date, no one has yet proposed a solution of the finite memory clearance time, since this will be at least as long as a delay itself! However, the higher the channel rate, the shorter the buffers need to be and the more practical the network becomes.

The other key function, again applicable both to WDM and OTDM-based packet networks, is address recognition. In electronic networks, addresses are loaded into a register and compared with a stored bit pattern. For very fast photonic networks, this is problematic. One proposed solution is to implement an optical AND gate (Cotter *et al.* 1997), which can route a packet according to its address, contained in an optical-routing header: a set of extra bits added to the packet. The synchronization pulse, derived from each packet, followed by the address of each packet, is compared with the keyword set for each node by a series of pre-fixed delays. The synchronization is key, since each node must be able to recognize the start of the packet, and assuming there is no network-wide clock (almost impossible to implement), some sort of asynchronous operation where timing, and, thus, packet alignment, and all processing must be carried out at the packet rate and not at the bit-level. If there is a match, the output of the AND gate triggers the switching action. Further delays/buffering must be introduced to avoid any contention. One rather ambitious proposal to avoid memory or buffering has been that of deflection, or 'hot-potato' routing, where in case of packet contention, the packet will be deflected and sent from node to node in a nomadic fashion until it can reach its destination. It is far from clear whether this is realistically implementable.

The clear perceived advantage with OTDM is the possible regeneration, demultiplexing and switching. All these functions can be realized using material nonlinearities either in semiconductors or fibre, based on interferometric switches whose refractive index is intensity dependent, as shown in figure 17. These can be made to operate as AND, NOT and XOR gates (see, for example, an excellent review in Chan *et al.* (1998)). Whether these will be implemented in semiconductors or as optical-fibre-based devices is unclear. While functionally the same, their time responses are different: semiconductors are compact but limited in speed to picosecond operation, while nonlinear fibre devices are relatively long (tens of metres) but have significantly faster switching speeds (sub-picosecond). To date, all-optical regeneration using an ultra-fast nonlinear interferometer has been shown at 40 Gbit s^{-1} (Thiele *et al.* 1999), demultiplexing to 10 Gbit s^{-1} in a nonlinear optical loop mirror from a myriad of rates up to 640 Gbit s^{-1}, and packet recognition at rates of up to 100 Gbit s^{-1} (Chan *et al.* 1998). The design and experimental demonstration of all necessary network functions out of these building blocks is one of the challenges for the future.

8. Future battles for bandwidth

The bandwidth in the optical domain, unlike almost any other part of the radio and microwave spectrum, is 'up for grabs'. There are numerous new optical devices, and an enormous transmission capacity. Yet the best way of using the bandwidth, whether through time-division or wavelength-division multiplexing, processing it as packets, or transparent high-capacity pipes, is far from clear.

Curiously, even researchers in the field could not have predicted the rate at which the bandwidth is now being consumed, nor how bandwidth hungry the global population has become. People and organizations appear to be capable of generating unimaginably large quantities of data. Perhaps the implication for society of this information expansion should be questioned. More information is only good for society if it is accurate, useful information, not waffle or junk email. Certainly, new information technologies are likely over the long run to alter patterns of employment, leisure and lifestyle across the world.

These, however, are mainly issues for governments, businesses and the general public. As we move into the next millennium, the research scientists and engineers working in this field face a big enough challenge: to continue to design new networks that carry the information as far and as fast as possible and can process it flexibly, quickly and reliably.

I thank The Royal Society for supporting me through a Royal Society University Research Fellowship and in many other ways. Huge thanks to Professor John Midwinter, and our past and present PhD students and research fellows in the Optical Networks Group, UCL, in particular, Stefano Baroni, Evgeny Churin, Michael Dueser, Paolo Gambini, Robert Killey, Vitaly Mikhailov, Derek Rothnie, Alex Stavdas, Hans Joerg Thiele, and many other colleagues, collaborators and funders, for valuable contributions to this paper and for enabling the active and exciting research work in this area. I am indebted to Sir Eric Ash and Robert Killey for critical reading of this paper.

References

Baroni, S. & Bayvel, P. 1997 Wavelength requirements in arbitrarily connected wavelength routed optical networks. *IEEE J. Lightwave Technol.* **15**, 242.

Baroni, S., Bayvel, P., Gibbens, R. J. & Korotky, S. 1999 Design of resilient multifibre wavelength-routed optical transport networks. *IEEE J. Lightwave Technol.* **17**, 743.

Bergano, N. & Davidson, C. R. 1995 Circulating loop transmission experiments for the study of long-haul transmission systems using erbium-doped fibre amplifiers. *IEEE J. Lightwave Technol.* **13**, 879.

Bergano, N. S. (and 14 others) 1995 100 Gbit/s WDM transmission of twenty 5 Gbit/s NRZ data channels over transoceanic distances using a gain flattened amplifier chain. In *Proc. 21st Eur. Conf. Optical Communications, Brussels, September*, vol. 3, pp. 967–970.

Bergano, N. S. (and 14 others) 1998 320 Gbit/s WDM Transmission (64×5 Gbit/s) over 7200 km using large mode fiber spans and chirped return-to-zero signals. In *Proc. OFC '98*, postdeadline paper PD-12.

Bergano, N. S., Davidson, C. R., Chen, C. J., Pedersen, B., Mills, M. A., Ramanujam, N., Kidorf, H. D., Puc, A. B., Levonas, M. D. & Abdelkader, H. 1999 640 Gbit/s transmission of 64×10 Gbit/s WDM channels over 7200 km with 0.33 nm (bit/s)/Hz spectral efficiency. In *Proc. OFC '99, San Diego*, postdeadline paper PD-2.

Chan, V. W. S., Hall, K. L., Modiani, E. & Rauschenbach, K. A. 1998 Architectures and technologies for high-speed optical data networks. *IEEE J. Lightwave Technol.* **16**, 2146.

Churin, E. & Bayvel, P. 1999 Free-space, dense WDM router based on a new concave grating configuration. *IEEE Photonics Technol. Lett.* **11**, 221–223.

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

Cotter, D., Lucek, J. K. & Marcenac, D. D. 1997 Ultra-high-bit-rate networking: from the transcontinental backbone to the desktop. In *IEEE Communications Magazine* (April), pp. 90–95.

Di Giovanni, D. J., Jablonowski, D. P. & Tan, M. F. 1997 Advances in fiber design and processing. In *Optical fiber telecommunications* (ed. I. P. Kaminow & T. L. Koch), vol. IIIA, ch. 4. Academic.

Gnauck, A. H. & Jopson, R. M. 1997 Dispersion compensation for optical fiber systems. In *Optical fiber telecommunications* (ed. I. P. Kaminow & T. L. Koch), vol. IIIA, ch. 7. Academic.

Hill, G. R. 1988 A wavelength-routing approach to optical communication networks. *BT Technol. J.* **6**, 24.

Hui, J. Y. & Cheung, K. W. 1996 Optical versus electronic switching for broadband networks. *IEEE Network* **10** (no. 6, special issue on Optical Networks), 21.

Hunter, D. K., Chia, M. C. & Andonovic, I. 1998 Buffering in optical packet switching. *IEEE J. Lightwave Technol.* **16**, 2081.

Hurm, V. 1998 40 Gbit/s, 1.55 μm monolithic integrated GaAs based PIN–HEMT photoreceiver. In *Proc. 24th Eur. Conf. Optical Communication (ECOC '98), Madrid, September*, vol. 3, pp. 119–121.

Imai, K., Tsuritani, T., Takeda, N., Tanaka, K., Edagawa, M. & Suzuki, M. 1999 500 Gbit/s (50×10 Gbit/s) WDM transmission over 4000 km using broadband EDFAs and low dispersion slope fibre. In *Proc. OFC '99*, postdeadline paper PD-5.

Kaiser, G. E. 1999 A review of WDM technology and its applications. *Opt. Fib. Technol.* **5**, 3.

Kaminow, I. P. & Koch, T. L. 1997 *Optical fiber telecommunications*, vols IIIA and IIIB. Academic.

Kawanishi, S. 1999 3 Tbit s^{-1} (160 Gbit/s × 19 channels) OTDM-WDM transmission experiments. In *Proc. OFC '99*, postdeadline paper PD-1.

Kawanishi, S., Takara, H., Uchiyama, K., Shake, I., Kamatani, O. & Takahashi, H. 1997 1.4 Tbit s^{-1} (200 Gbit/s × 7 channels) 50 km optical transmission experiment. *Electron. Lett.* **33**, 1716.

Keshav, S. & Sharma, R. 1998 Issues and trends in router design. In *IEEE Commun. Mag. (May 1998)*, pp. 144–151.

Koch, T. L. 1997 Laser sources for amplified and WDM lightwave systems. In *Optical fiber telecommunications* (ed. I. P. Kaminow & T. L. Koch), vol. IIIB, ch. 4. Academic.

Kumar, V. P., Lakshman, T. V. & Stiliadis, D. 1998 Beyond best effort: router architectures for the differentiated services of tomorrow's internet. In *IEEE Commun. Mag. (May 1998)*, pp. 152–164.

Lagasse, P. (and 10 others) 1998 *Photonic technologies in Europe* (ed. C. Descamps). Telenor. (Thematic issues from ACTS: http://www.inforwin.org.)

Li, P. Y. & Henry, C. H. 1997 Silicon optical bench wave guide technology. In *Optical fiber telecommunications* (ed. I. P. Kaminow & T. L. Koch), vol. IIIB, ch. 8. Academic.

Lin, L. Y., Lunardi, L. M. & Goldstein, E. L. 1998 Optical cross-connect integrated systems: a free-space micromachined module for signal and switching configuration monitoring. In *Proc. IEEE/LEOS Summer Topical Meeting on 'Optical MEMs'*, paper PD004.

Malyon, D., Widdowson, T., Bryant, E. G., Carter, S. F., Wright, J. V. & Stallard, W. A. 1991 Demonstration of optical pulse propagtion over 10 000 km using a recirculating loop. *Electron. Lett.* **27**, 120.

Marcuse, D. & Menyuk, C. 1999 Simulation of single channel optical systems at 100 Gbit/s. *IEEE J. Lightwave Technol.* **17**, 564.

Morita, I. 1998 1.04 Tbit s^{-1} DWDM transmission experiment based on alternate polarisation 80 Gbit/s OTDM signals. In *Proc. 24th Eur. Conf. Optical Communications* (postdeadline papers), vol. 3, pp. 53–55.

Murakami, M. 1998 Quarter terabit (25 × 10 Gbit/s) over 9288 km WDM transmission experiment using nonlinear supported RZ pulse in higher order fiber dispersion managed line. In *Proc. 24th Eur. Conf. Optical Communications, Madrid, September* (postdeadline paper), vol. 3, pp. 77–79.

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.

Sabry, M. & Midwinter, J. E. 1994 Towards the optical ether. In *IEE Proc. Optoelectron.* **141**, 327.

Taga, H. 1998 213 Gbit/s (20 × 10.66 Gbit/s) over 9000 km transmission experiment using a dispersion slope compensator. In *Proc. OFC '98*, postdeadline paper PD13.

Takahata, K., Miyamoto, Y., Muramoto, Y., Fukano, H. & Matsuoka, Y. 1998 50 Gbit/s operation of a monolithic wave guide pin photodiode/HEMT (high-mobility electron transistors) receiver opto-electronic IC module. In *Proc. 24th Eur. Conf. Optical Communication (ECOC '98), Madrid, September*, vol. 3, pp. 67–69.

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

Yamada, M., Mori, A., Kobayashi, K., Ono, H., Kanamori, T., Oikawa, K., Nishida, Y. & Ohishi, Y. 1998 Gain flattened tellurite-based EDFA with a flat amplification bandwidth of 76 nm. In *Proc. OFC '98* postdeadline paper PD-7.

Yano, Y., Ono, T., Fukuchi, K., Ito, T., Yamazaki, H., Yamaguchi, M. & Emura, K. 1996 2.6 Tbit s^{-1} transmission experiment using optical duobinary coding. In *Proc. 22nd Eur. Conf. Optical Communications, Oslo, Norway* (postdeadline papers), vol. 5, pp. 3–6.

Yonenaga, K., Matsuura, A., Kuwahara, S., Yoneyama, M., Miyamoto, Y., Hagimoto, K. & Noguchi, K. 1998 Dispersion-compensation-free 40 Gbit $s^{-1} \times 4$ channel WDM transmission experiment using zero-dispersion-flattened transmission line. In *Proc. Optical Fiber Conf., OFC'98, Dallas*, paper PD20.

AUTHOR PROFILE

P. Bayvel

Polina Bayvel, 33, was born in the former USSR. She received her BSc(Eng.) and PhD degrees in Electronic and Electrical Engineering from University College London (UCL) in 1986 and 1990, respectively. In her PhD work she specialized in nonlinear fibre optics and its applications. Following a Royal Society postdoctoral exchange fellowship in the Fiber Optics Laboratory at the General Physics Institute in Moscow (USSR Academy of Sciences) in 1990, she worked as a Principal Systems Engineer at STC Submarine Systems Ltd (Greenwich, UK) and BNR Europe Ltd (subsequently Nortel Networks, Harlow, UK, and Ottawa, Canada) on the design and planning of high-speed optical transmission networks. In 1993 she was awarded a Royal Society University Research Fellowship at UCL. Since 1994 she has built up the Optical Networks research group whose research activities are in the areas of high-speed optical communication systems, networks and associated devices. She has authored/co-authored over 75 refereed journal and conference papers and serves on the Editorial Board of *Optical Fiber Technology*, on the organizing committee of the IEEE Lasers and Electro-Optics Society (LEOS) Annual Meeting, and the Royal Society steering committee on ‘Science and Society’.