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

Advances in Nonlinear Optics for Information Processing and All-Optical Networking

K. Smith, J. K. Lucek, R. J. Manning and K. J. Blow

Phil. Trans. R. Soc. Lond. A 1996 **354**, 707-717
doi: 10.1098/rsta.1996.0025

Advances in nonlinear optics for information processing and all-optical networking

By K. Smith, J. K. Lucek, R. J. Manning and K. J. Blow
BT Laboratories, Martlesham Heath, Ipswich IP5 7RE, UK

Telecommunications networks of the future will be expected to deal with a far greater range of bandwidth-demanding services than the current limited requirements of telephony and data traffic. In order to satisfy new customer demands, network designers are likely to consider novel techniques for switching and transmission in order to fully exploit the enormous bandwidth potential of optical fibre. In this paper, we will review recent advances in the application of nonlinear optical information processing techniques to all-optical networking.

1. Introduction

Optical fibre has become the preferred transmission medium for the modern day telecommunications network. In general, this network is characterized by a hybrid architecture, whereby high-bandwidth single-mode optical fibre links are used to interconnect electronic nodes. Switching and routing is performed entirely in the electrical domain owing to the advanced state of electronic switching technology. Such a network has evolved primarily out of the need to satisfy the limited requirements of (circuit-switched) telephony and (packet-switched) data traffic. The future is likely to be very different. Networks will be expected to deal with a far greater range of bandwidth-demanding services in addition to the traditional voice and data traffic. At the top end of the bandwidth scale are the multimedia applications with a mixture of traffic types including full motion video. In order to satisfy these new customer demands, network designers are likely to consider novel techniques for switching and transport in order to fully utilize the enormous bandwidth potential of an optical fibre infrastructure. It is very likely that the rich field of nonlinear optics can provide a number of revolutionary technological solutions. For example, in the case of optical solitons, the effect of nonlinearity in fibre optic links can be used to offset the chromatic dispersion in the fibre and thereby give rise to pulses which do not change shape or duration during propagation. Such nonlinear phenomena were once in the realms of esoteric science, but within the past 10–15 years have risen to become serious contenders in the race for the next generation of high-capacity long-haul transmission systems. Another major area of potential impact is in the use of optical switching technology. Here, nonlinear optical media can be exploited to allow beams of light to interact, i.e. we use light to control light. This new class of photonic switches are inherently fast (subpicosecond switching speeds are possible) and have high transmission bandwidths. Therefore, the relatively slow transmission and processing speeds of present day electronic devices could be overcome by en-

Phil. Trans. R. Soc. Lond. A (1996) **354**, 707–717
Printed in Great Britain

suring that the information is kept in an optical format throughout the switching operation.

In this paper, we will review the pertinent principles of nonlinear optics which underpin the description of soliton phenomena and light controlling light in fibres. To date, most of the key optical processing demonstrations have used fibre-based devices and we will describe the technology together with recent advances in functionality, such as signal regeneration and programmable routing nodes. We will also highlight a very recent major step forward in nonlinear material research, whereby nonlinear effects in semiconductors have been shown to be applicable to ultrahigh-speed ($100\ \mathrm{Gbit\ s^{-1}}$) processing applications.

2. Pertinent nonlinear optics

The first requirement for all-optical switching is a nonlinear medium in order for the switching and the switched beams to interact with each other. Our discussion in this section will be limited to the case of silica-based fibres where third-order nonlinearity (the optical Kerr effect) is used. The nonlinear phase shift, $\Delta\Phi_{\mathrm{signal}}$, incurred by the signal is therefore

$$\Delta\Phi_{\mathrm{signal}} = k_0 n_2 L I_{\mathrm{signal}} + 2k_0 n_2 L I_{\mathrm{switch}},$$

where n_2 is the Kerr coefficient ($\mathrm{m^2\ W^{-1}}$), k_0 is the propagation constant ($\mathrm{m^{-1}}$), L is the fibre length (m) and I is the beam intensity ($\mathrm{W\ m^{-2}}$). The first term (giving rise to self-phase modulation (SPM)) is the key to understanding soliton phenomena. Essentially, the optically induced frequency chirp via the SPM term can be used to offset the chirp incurred via chromatic dispersion in the fibre. This SPM term can also be used to achieve intensity-dependent routing and soliton switching (Blow *et al.* 1989). It is the second term that is of prime interest for allowing one (control) beam to switch another (signal). This is often referred to as the cross-phase modulation (XPM) term. One very useful feature of XPM is that we can tailor the temporal shape of the induced phase profile by exploiting the fact that the signal and switch pulses can travel at different speeds in the fibre. This effect modifies the imposed phase shift as follows:

$$\Delta\Phi_{\mathrm{signal}} = 2k_0 n_2 \int_0^L I_{\mathrm{switch}}(t - \Delta\beta' z)\,\mathrm{d}t,$$

where $\Delta\beta'$ is the group delay difference between pump and signal. This is illustrated in figure 1*b* for the cases of small (i) and large (ii) group delay difference or 'walk-off'. The width of the phase window can be tuned for a particular application. Experimentally, we achieve this by selection of appropriate linear fibre dispersion properties and wavelengths for the signal and control beams. In general, our fibre devices are operated such that the switching and signal wavelengths are chosen to straddle the dispersion minimum of the nonlinear fibre. Typically, a standard set up is to use a length of commercially available dispersion-shifted fibre ($\lambda_0 \sim 1.55\ \mu\mathrm{m}$) with nominal control and signal wavelengths of 1.54 and 1.56 μm, respectively. Owing to the extremely 'small' value of n_2 in silica fibres ($n_2 \sim 3.2 \times 10^{-20}\ \mathrm{m^2\ W^{-1}}$), long ($\sim$ km) lengths of fibre are usually required in order to achieve full switching. For the case of no walk-off, the peak power required for a phase change of π is given by $P_\pi L = \lambda A/4n_2$. For $\lambda = 1.55\ \mu\mathrm{m}$ and an effective core area, $A = 50\ \mu\mathrm{m}^2$, we obtain $P_\pi L = 0.6$ W km, which means that for a 1 km long loop, we require ~ 0.6 W

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

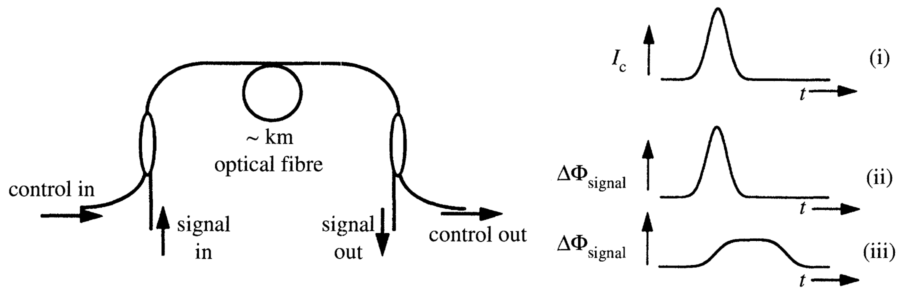

(*a*) (*b*)

Figure 1. (*a*) Control and signal beams are combined in a length of optical fibre resulting in the nonlinear profiles depicted in (*b*). The phase profile $\Delta\Phi_{\mathrm{signal}}$ corresponding to the control pulse ((i)) are shown in (ii) and (iii) for the cases of small and large 'walk-off', respectively.

peak power for a π phase change. (For a 100 Gbit s^{-1} system, with pulse durations $\sim$ 1 ps, a mean power of $\sim$ 30 mW would be required.)

3. Optical solitons

Optical solitons may prove to be the fundamental 'bits' in high-speed networks. A full discussion of the properties of solitons is beyond the scope of this paper, but can be found elsewhere (Blow & Doran 1990). A strict analysis of the phenomenon requires solving the wave equation in the presence of nonlinearity (SPM) and dispersion. The resulting equation once normalized is referred to as the nonlinear Schrödinger equation (NLSE). Although the NLSE has an infinite number of solutions, the most practically significant solution so far is the so-called fundamental soliton, the width and shape of which do not change during propagation. In this case, the two effects (nonlinearity and dispersion) have equal magnitude with opposite signs and therefore cancel each other out. For silica fibres, this means that the dispersion must be anomalous (i.e. the signal wavelength is greater than the wavelength of minimum dispersion). Also, the amplitudes of the soliton solutions are not arbitrary—they are determined by the magnitudes of the nonlinearity and the dispersion.

The above description of soliton phenomena is somewhat idealized. In real fibres, the soliton is subjected to a number of potentially deleterious effects. The fibre loss is particularly detrimental since as the power decreases during propagation, the nonlinear contribution weakens and the soliton eventually broadens due to the dispersion. In order to offset this effect, the use of optical amplifiers was proposed (Hasagawa 1983). The first experimental demonstration illustrating soliton transmission over many thousands of kilometres using optical amplifiers was performed using periodic loss compensation via Raman gain in a recirculating loop configuration (Mollenauer & Smith 1988). Although these pioneering experiments illustrated the true potential of soliton systems, it was not until the advent of practical fibre amplifiers (in particular, the erbium-doped fibre amplifier) that solitons really came of age.

Soliton phenomena have now been exploited in major laboratories all over the world in order to send optical data at high speeds and over long distances (Widdowson & Ellis 1994; Mollenauer *et al.* 1993; Nakazawa *et al.* 1993). There are still limitations, however, which arise primarily from the amplified spontaneous emission (ASE) noise inherent in the optical amplification process. As the solitons propagate along a chain of amplifiers, the vagaries of ASE lead to both a degradation in the signal to noise level and to a random pulse-to-pulse timing jitter (Gordon & Haus 1986). Nonlinear phenomena may be used to circumvent these limitations. An elegant approach to the problem is the use of 'sliding-frequency guiding' filters (Mollenauer *et al.* 1992), which rely on the fundamentally nonlinear nature of the solitons. Briefly, in this passive technique, the centre wavelengths of the optical filters are translated (by a few pulse bandwidths) over the total span of the amplifier chain. The soliton, owing to its nonlinear nature, is able to generate new frequency components and thereby follow the translation imposed by the filters. The ASE, however, having no 'line of sight' through the filter ensemble, is attenuated, therefore allowing preferential transmission of the soliton. Another attractive passive technique is to use intensity-dependent switching (Blow *et al.* 1989; Smith *et al.* 1992), which acts as an intensity discriminator (or soliton filter), transmitting the soliton in preference to the low-level ASE noise. In the following section, we will describe solutions involving active optical processing techniques which not only enhance the capacity of the all-optical links but are also used to perform key processing functions at the heart of the all-optical network nodes.

4. Optical processing—building blocks and functionality

In this section, we look at optical processing building blocks and show how novel and simple optical processing schemes can be used to enhance (soliton) transmission capacity and how data can be demultiplexed and routed in an all-optical network. The optical circuits described in this section are based on the Kerr effect—the process of XPM discussed previously. The simple building block of the devices in this work is therefore the all-optical fibre phase modulator, illustrated in figure 1*a*. In order to perform optical switching, we use some form of interferometer. This is commonly achieved using a Sagnac loop configuration, as shown in figure 2*a*, which since its first demonstration (Blow *et al.* 1990) has become the benchmark optical switching device (Nelson *et al.* 1991; Andrekson *et al.* 1992; Uchiyama *et al.* 1994). The two paths in the interferometer are the clockwise and anti-clockwise propagating signals (at λ_s). In the absence of a control beam (at λ_c), we can ensure that the counter-propagating signals interfering at the loop coupler are either fully reflected or transmitted depending on the adjustment of the polarization controllers (PC) in the loop. Optical switching is typically achieved by introducing the control beam using wavelength dependent couplers such that it travels in only one direction around the loop (clockwise in figure 2*a*). If the control pulse overlaps the signal pulse then the phase of the clockwise signal is altered via XPM. For a 50 : 50 loop coupler, the transmission of the loop is given by

$$P_{\text{out}} = P_{\text{in}} \tfrac{1}{2}\{1 - \cos(\Delta\Phi)\},$$

where P_{out} and P_{in} are the output and input powers, respectively, and $\Delta\Phi$ is the phase difference between the two counter-propagating beams. As we saw from figure 1, the shape of the imposed phase profile depends on the amount of 'walk-off'. We therefore have a simple way of ensuring a 'square' switching window with an adjustable 'on' time. This is particularly useful in suppressing the deleterious effects of timing jitter, for example, when using optical data to switch a clock pulse stream (e.g. optical signal regeneration).

The above all-optical AND gate device is often referred to as a nonlinear optical loop mirror (NOLM) and has been one of the key building blocks in the demonstration

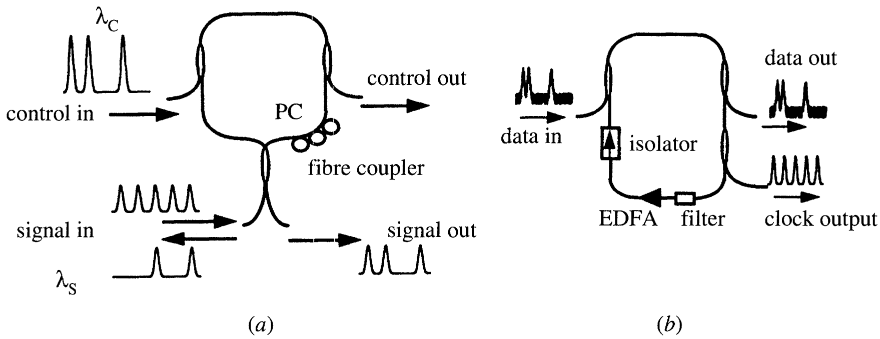

Figure 2. Two of the all-fibre optical processing building blocks: (*a*) the nonlinear optical loop mirror (NOLM); (*b*) the optically mode-locked (clock recovery) laser based on an erbium-doped fibre amplifier (EDFA).

of more advanced optical processing functionality. Another requirement of all-optical networks is that of synchronization. In the context of a synchronous network architecture, this amounts to the question: how do we generate a continuous stream of optical pulses which are locked to the phase and frequency of the incoming data? In our work we have used the technique (Smith & Lucek 1992) depicted in figure 2*b*, which is based on a mode-locked laser. Since a laser cavity and a transmission fibre share the same optical modulator (which in our demonstration was the phase modulator fibre in figure 1*a*), the optical data can serve to internally modulate the light in the laser cavity. Providing that this modulation takes place with a period (T) equal to (or an integer multiple of) the laser round-trip time, mode-locking of the laser follows. Since the mode-locked laser is essentially an injection locked oscillator, the ensuing continuous stream of high-quality ps-duration pulses (at a repetition rate of $1/T$) is now precisely locked to the data stream and, therefore, any low-frequency phase shifts on the incoming data can be tracked by the laser. This network synchronization function or 'clock recovery' as it is called, is an essential ingredient in any remote processing architecture and has been demonstrated at speeds up to 40 GHz (Ellis *et al.* 1993).

We will now focus on examples of the more advanced optical processing functionality that has been demonstrated using the building blocks described above. The first is illustrated in figure 3*a*, which shows a schematic of an all-optical signal regenerator (Lucek & Smith 1993). The first stage is the all-optical clock recovery circuit. In this way, we derive a continuous stream of optical pulses which are in precise synchronism with the incoming data. In the next stage of regeneration, the data exiting the clock recovery stage is used a second time to amplitude modulate the clock train (using a NOLM), thus the clock pulses become the regenerated data stream. Since the clock pulses emerge both reshaped and with reduced timing jitter, true signal regeneration is achieved. Results from our first experiment are shown in figure 4. In this example, walk-off between control and signal in the NOLM is particularly important in ensuring that the temporal jitter on the data does not lead to undesirable (amplitude) noise on the modulated clock.

A particularly elegant method of enhancing the transmission performance of soliton networks is the technique that we refer to as 'soliton shepherding' (Widdowson *et al.* 1994). The basic concept is illustrated in figure 5*a*: by temporally overlapping a clock and (soliton) data stream in a nonlinear optical medium (NOM), we arrange for a nonlinear phase profile (via XPM) to be imposed on the solitons by the clock pulses.

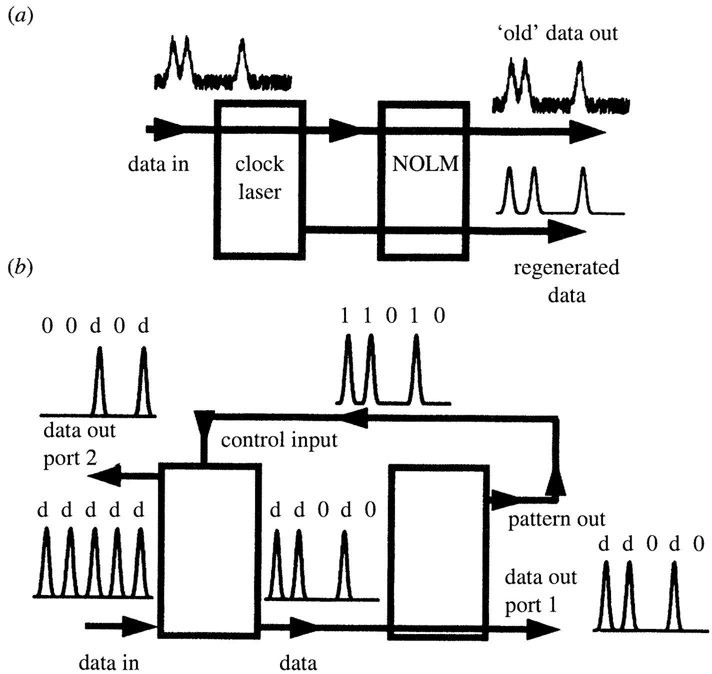

Figure 3. Schematic diagrams of (*a*) the all-optical signal regeneration and (*b*) the remotely programmable network node.

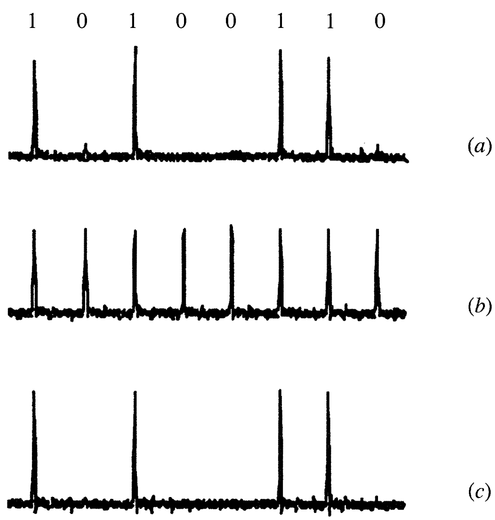

Figure 4. Experimental results of all-optical signal regeneration: (*a*) the 8-bit input pattern (10100110) to the regenerator (at 1.54 μm); (*b*) the recovered clock (at 1.56 μm); (*c*) the regenerated data.

Owing to the particle-like nature of the soliton, the imposed phase modulation is distributed over the whole pulse, resulting in a net shift in its carrier frequency. The sign and the magnitude of the frequency shift depend on the position of the soliton relative to the clock pulse (figure 5*b*). If we combine this effect with an appropriate sign and magnitude of fibre dispersion, the imposed frequency shifts can be converted

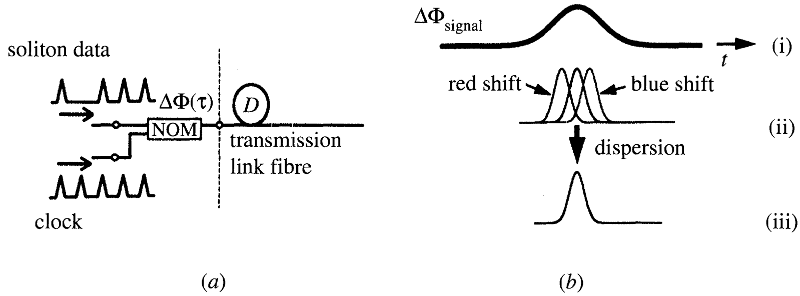

Figure 5. (*a*) The soliton shepherding concept and (*b*) the nonlinear phase profile (shown in (i)) gives rise to temporally dependent shifts in the optical frequency (shown in (ii)), which, under the influence of anomalous dispersion, are converted into appropriate time delays (iii).

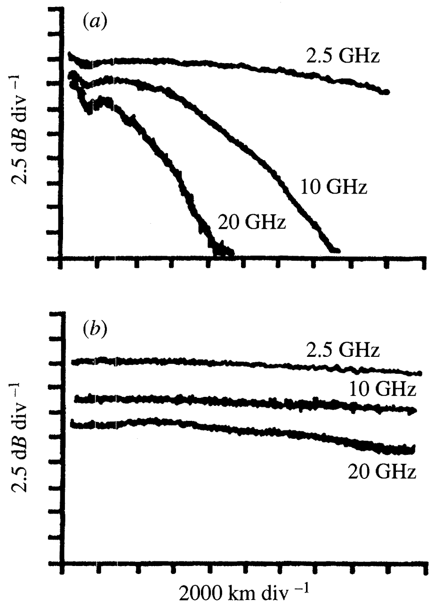

Figure 6. Microwave components (2.5, 10 and 20 GHz) of 2.5 Gbit s^{-1} soliton data as a function of distance for (*a*) no shepherding beam, and (*b*) with weak shepherding beam.

into suitable time delays. Hence, the clock acts to 'shepherd' the soliton towards the centre of the time slot defined by the clock pulse. In the case demonstrated here, the NOM is a simple length of optical fibre. In fact, the fibre is actually the same fibre that is used in the transmission link. The XPM together with the soliton supporting dispersion present within the transmission link act to temporally guide the soliton bits. Preliminary experiments have shown that only a weak shepherding beam (giving rise to peak phase shifts of only ~ 0.01 rad) is sufficient to suppress timing jitter, as illustrated in the experimentally recorded microwave traces depicted in figure 6. In this example, the error-free (10^{-9}) distance was extended significantly from 9000 km to more than 18 000 km using the soliton shepherding technique.

In principle, the above soliton regeneration and retiming techniques offer the potential of sending high-speed soliton data over almost unlimited distances. A problem we must then address is how to manipulate this high-speed data to perform demultiplexing and routing. A solution is illustrated in figure 3*b*. Here we have the same

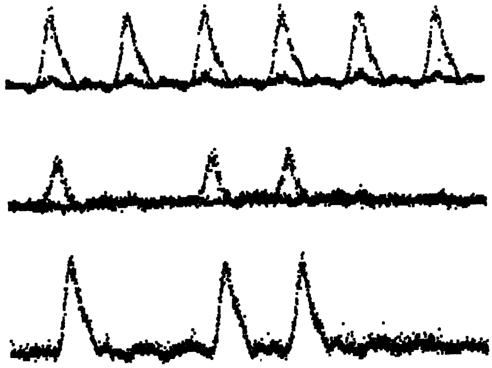

Figure 7. Experimental results (output of 2.5 GHz receivers displayed on a sampling oscilloscope triggered at the frame rate): (*a*) incoming data; (*b*) data routed to port 1 (figure 3*b*); (*c*) recovered optical clock pattern.

building blocks as in the regenerator, but the elements are 'wired' in a different way. The new configuration allows us to demultiplex (or route) an arbitrary combination of channels from an incoming optical time division multiplexed (OTDM) data stream (Lucek & Smith 1995). The configuration comprises a NOLM and an optically mode-locked laser clock recovery stage operating in a novel mode that we call optical clock pattern recovery. The output of the clock circuit is then fed back to control the NOLM. Any arbitrary combination of incoming channels can be demultiplexed (or routed). Results are shown in figure 7, where we have chosen to demultiplex the first, third and fourth channels of an incoming six-channel OTDM stream. The device is remotely programmable because the routing of a channel depends on whether the data source was transmitting data in that channel during the start-up of the device—there is nothing in the hard 'wiring' that determines the port to which a channel is routed. The clock pattern recovery property of the mode-locked laser gives the routing device the flexibility to route any combination of channels to a given port. In the most general sense, the device is performing a divided frequency optical clock generation (or optical prescaling) function. This divided optical clock can be employed in further processing functions as well as gating the NOLM.

5. Advances in nonlinear materials

The above work has illustrated some of the functionality which has now been demonstrated using all-optical processing technology. Over the past year or so, there have been some startling experiments illustrating the applicability of semiconductors in performing ultrahigh-speed optical processing. The first experiment by Ellis & Spirit (1993) used a semiconductor laser amplifier (SLA) in a fibre loop mirror configuration to demultiplex a 10 Gbit s^{-1} channel from a 40 Gbit s^{-1} OTDM data stream. Latterly, a similar configuration has been used to demultiplex a 10 Gbit s^{-1} channel from a line rate of 160 Gbit s^{-1} (Suzuki *et al.* 1994). The SLA is the nonlinear element in a loop mirror and is placed slightly offset from the centre of the loop fibre path (as shown in figure 8). Data injected into the loop is split by a 50:50 coupler and the two streams arrive asynchronously at the SLA. The SLA is addressed by 10 GHz gating pulses, which arrive after one data pulse but before its counter-propagating replica. This imparts a π relative phase shift on to the replica, so the pulse is switched out at the coupler.

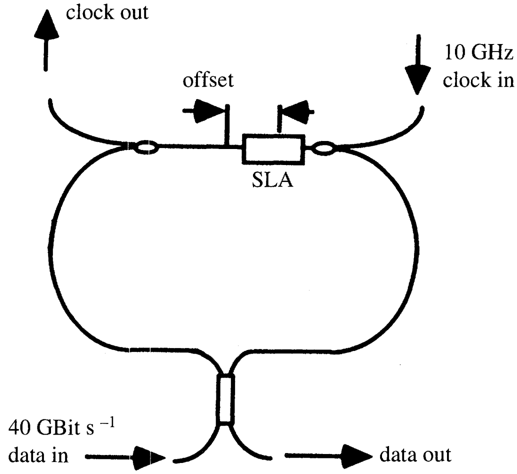

Figure 8. The SLA-loop mirror configuration used in the high-speed demultiplexing experiment of Ellis & Spirit (1993).

The fact that the switching could take place at a repetition rate of 10 GHz, much faster than the recovery rate of the SLA (1 GHz), was very surprising, given previous understanding (Sokoloff *et al.* 1993; Eiselt *et al.* 1993). As such, it was the first of several key steps towards a change of view of the conventional switching requirements for materials, which had assumed that slow resonant nonlinearities were inappropriate for ultrafast high repetition rate switching. Basically, a π phase recovery is only a fraction of the total recovery when the gain of the SLA is run into quasi-saturation by a repetitive gating pulse. The SLA, therefore, never recovers between gating pulses, but does not need to, since, having recovered by π, it is ready to switch again.

The second step towards the successful exploitation of SLAs in the circuitry described previously was its incorporation in a clock recovery circuit as the nonlinear element (Patrick & Manning 1994). Here an intense holding beam was required to speed the response of the SLA so as to prevent phase and amplitude wander in the presence of zeros in the incoming data stream. This technique successfully demonstrated clock recovery at 20 Gbit s^{-1}. It should be noted that, in the context of wavelength division multiplexing applications, wavelength conversion at 20 Gbit s^{-1} has been demonstrated using a saturating beam to both speed the response and sample the data (Stubkaer *et al.* 1994).

More recently, the key regenerative function of data switching using a SLA in a loop mirror has been demonstrated at 10 Gbit s^{-1} (Davies *et al.* 1995). Here, data switches pulses from a regular clock stream in the loop. Although data switching means that phase wander occurs in the presence of zeros, the topology of the loop mirror arrangement and its transmission characteristic mean that error-free switching is possible. The most recent results on phase recovery speed in SLAs have shown that π recovery is possible in times as short as 12.5 ps (Manning & Sherlock 1994). This result (figure 9), based on an interferometric measurement technique, indicates that it should be possible to switch at a gating rate of 80 GHz. Ultrafast (fs) measurements (Hall *et al.* 1994) show that transients last up to a few ps, so it is likely that switching at rates of 100 GHz and beyond should be achievable. These advances mean that a version of the high-speed all-optical regenerator made mostly or entirely out of SLAs

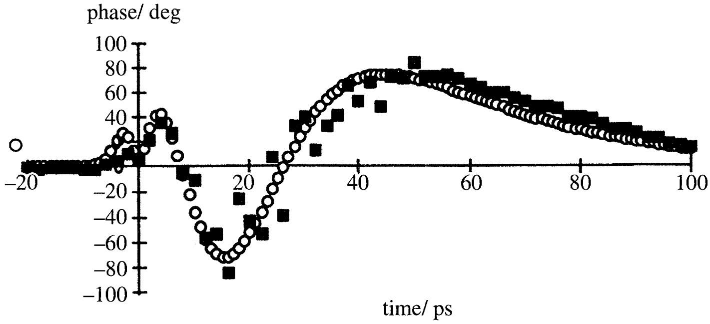

Figure 9. Time resolved fringe shift measurement in a SLA. The measured phase swings between $\pm\frac{1}{2}\pi$ due to the experimental measurement technique (Manning & Sherlock 1994). The first 0.8π recovers in 10 ps. Open circles are theoretical points (calculated assuming a 27 ps lifetime and 4 ps pulse duration), full squares are experimental data.

is possible, giving an enormous reduction in the size of these devices, and allowing the possibility of further integration.

6. Conclusions

In this paper, we have reviewed some of the areas where the field of nonlinear optics could play a key role in defining new technologies for future high-speed all-optical telecommunications networks. The past few years have seen major advances in the levels of optical processing functionality with pioneering demonstrations of novel signal regeneration and routing techniques. Using optical regeneration techniques, we have the possibility of sending high-speed (100 Gbit s^{-1} and beyond) soliton data over almost unlimited distances. In addition, by keeping data in an optical format throughout the switching process, optical routing offers an elegant and simple solution to the manipulation of these high-speed data streams and could eventually replace some of the present electronic devices. It is also highly likely that by exploiting the recent advances in the understanding of semiconductor-based nonlinearities and devices, the above processing functions could be performed using mostly or entirely semiconductor materials. In this way, device integration may be achieved. This opens up new and exciting opportunities in the domain of very high functionality ultrafast processing. Here, optical technology may even penetrate and determine the design of tomorrows computing structures.

References

Andrekson, P. A., Olsson, N. A., Simpson, J. R., Digiovanni, D. J., Morton, P. A., Tanbun-Ek, T., Logan, R. A. & Wecht, K. W. 1992 Ultrahigh speed demultiplexing with the nonlinear loop mirror. Technical Digest of OFC'92, PD8, pp. 343–346.

Blow, K. J. & Doran, N. J. 1990 Solitons in optical fibres. In *Nonlinear waves in solid state physics* (ed. A. D. Boardman). New York: Plenum.

Blow, K. J., Doran, N. J. & Nayar, B. K. 1989 Experimental demonstration of optical soliton switching in an all-fibre nonlinear Sagnac interferometer. *Opt. Lett.* **14**, 754–756.

Blow, K. J., Doran, N. J., Nayar, B. K. & Nelson, B. P. 1990 Two-wavelength operation of a nonlinear fibre loop mirror. *Opt. Lett.* **15**, 248–250.

Davies, D. A. O., Ellis, A. D., Widdowson, T. & Sherlock, G. 1995 10 Gbit s^{-1} data switched semiconductor laser amplifier nonlinear loop mirror. *Electron. Lett.* (To be published.)

Ellis, A. D., K. Smith & Patrick, D. M. 1993 All-optical clock recovery at bit-rates up to 40 Gbit s^{-1}. *Electron. Lett.* **29**, 1323–1324.

Ellis, A. D. & Spirit, D. M. 1993 Compact 40 Gbit s^{-1} optical demultiplexer using a GaInAsP optical amplifier. *Electron. Lett.* **29**, 2115–2116.

Eiselt, M., Pieper & Weber 1993 All-optical high-speed demultiplexing with a semiconductor laser amplifier in a loop mirror configuration. *Electron. Lett.* **29**, 1167–1168.

Gordon, J. P. & Haus, H. A. 1986 Random walk of coherently amplified solitons in optical fibre transmission. *Opt. Lett.* **11**, 665–667.

Hall, K. L., Lenz, G., Darwish, A. M. & Ippen, E. P 1994 Subpicosecond gain and index nonlinearities in InGaAsP diode lasers. *Opt. Commun.* **111**, 589–612.

Hasagawa, A. 1983 Amplification and reshaping of optical solitons in glass fibres IV. *Opt. Lett.* **8**, 650–652.

Lucek, J. K. & Smith, K. 1993 All-optical signal regeneration. *Opt. Lett.* **18**, 1226–1228.

Lucek, J. K. & Smith, K. 1995 Remotely programmable routing device with optical clock division. *IEEE Photonic. Tech. Lett.* **7**, 59–61.

Manning, R. J. & Sherlock G. 1994 π Recovery on $\sim$ 10 ps timescale in a semiconductor laser amplifier. Postdeadline paper PDP-30 at OSA annual meeting, Dallas.

Mollenauer, L. F. & Smith, K. 1988 Demonstration of soliton transmission over more than 4000 km in fibre with loss periodically compensated by Raman gain. *Opt. Lett.* **13**, 675–677.

Mollenauer, L. F., Gordon, J. P. & Evangelides, S. G. 1992 The sliding-frequency guiding filter: an improved form of soliton jitter control. *Opt. Lett.* **17**, 1575–1577.

Mollenauer, L. F., Lichtman, E., Neubelt, M. J. & Harvey G. T. 1993 Demonstration, using sliding-frequency guiding filters, of error-free soliton transmission over more than 20 Mm at 10 Gbit s^{-1}, single channel, and over more than 13 Mm at 20 Gbit s^{-1} in a two channel WDM. *Electron. Lett.* **29**, 910–912.

Nakazawa, M., Suzuki, K., Yodama, E., Kimura, Y. & Toboya, M. 1993 Experimental demonstration of soliton data transmission over unlimited distances with soliton control in time and frequency domain. *Electron. Lett.* **29**, 729–730.

Nelson, B. P., Blow, K. J., Constantine, P. D., Doran, N. J., Lucek, J. K., Marshall, I. W. & Smith, K. 1991 All-optical Gbit s^{-1} switching using a nonlinear loop mirror. *Electron. Lett.* **27**, 704–705.

Patrick, D. M. & Manning, R. J. 1994 20 Gbit s^{-1} all-optical clock recovery using semiconductor nonlinearity. *Electron. Lett.* **30**, 151–152.

Smith, K., Greer, E. J., Doran, N. J., Bird, D. M. & Cameron, K. H. 1992 Pulse amplification and shaping using a nonlinear loop mirror that incorporates a saturable gain. *Opt. Lett.* **17**, 408–410.

Smith, K. & Lucek, J. K. 1992 All-optical clock recovery using a mode-locked laser. *Electron. Lett.* **28**, 1814–1815.

Sokoloff, J. P., Prucnal, P. R., Glesk, I. & Kane, M. 1993 A terahertz optical asymmetric demultiplexer (TOAD). *IEEE Photonics Tech. Lett.* **5**, 787–790.

Stubkjaer, K. E. *et al.* 1994 Optical wavelength convertors. In *ECOC'94 Proc.*, vol. 2, pp. 635-642.

Suzuki, K., Iwatsuki, K., Nishi, S. & Saruwatari, M. 1994 160 Gbit s^{-1} single polarisation subpicosecond transform limited pulse signal demultiplexing using ultrafast optical loop mirror including MQW travelling wave semiconductor laser. *Electron. Lett.* **30**, 660–661.

Uchiyama, K., Kawanishi, S., Takara, H., Morioka, T. & Saruwatari, M. 1994 100 Gbit s^{-1} to 6.3 Gbit s^{-1} demultiplexing experiment using a polarisation-independent nonlinear optical loop mirror. *Electron. Lett.* **30**, 873–874.

Widdowson, T., Malyon, D. J., Ellis, A. D., Smith, K. & Blow, K. J. 1994 Soliton shepherding: all-optical active soliton control over global distances. *Electron. Lett.* **30**, 990–991.

Widdowson, T. & Ellis, A. D. 1994 20 Gbit s^{-1} soliton transmission over 125 Mm. *Electron. Lett.* **30**, 1866–1867.