All-optical regeneration will rewrite the rulebook

GLOBAL CROSSING

After wavelength-division multiplexing and optical switching, the next likely dislocation in photonic technology is already taking shape. John Devaney talked to two of the researchers currently setting the agenda on all-optical signal regeneration.

WHAT EVERY telecoms entrepreneur wants to find, apart from a generous venture capitalist, is a technical wizard with a "dislocating" technology. That is, an innovation that opens up a new revenue stream for the industry, giving the entrepreneur and backers first-mover advantage.

Wavelength-division multiplexing was one such dislocation, with US equipment maker Ciena as one of its main beneficiaries. The market conditions, the maturity of the technology and, possibly most important of all, the price came together to create a compelling market opportunity at just the right time.

Another convergence of circumstances is likely to occur as optical line rates climb – to 40 Gbit/s and beyond – and optical signal routing becomes commonplace across the network. In this scenario, high-bit-rate wavelength signals will become more difficult to transmit over significant distances. The counterpoint to this, of course, is the fact that carriers want hassle-free global networking, with built-in flexibility, scalability and rapid provisioning all in the optical domain.

At the moment, when bit-error rates become dangerously high, every channel is demultiplexed, converted into the electrical domain and given a health check before the optical signal is recreated, remultiplexed and retransmitted down the next leg of its journey. Backbone networks are dotted with these optical–electrical–optical (OEO) regenerators, which eat up large quantities of space, power and capital investment.

The alternative is to do the signal regeneration in the optical domain, something that has been achieved on the laboratory bench, but that is not yet commercially viable. The technology is likely to be expensive, but then so too are the banks of optical–electrical transponders that make up today's OEO regenerators. What's more, those transponders are going to get a lot more costly as line rates get faster.

Although the optical regeneration techniques currently under development effectively use optical processing to recreate the original signal, regeneration is traditionally broken down into three components: reamplification, retiming and reshaping. Hence the term "3R regeneration".

Rocky road ahead

Today, the best fibre available has just less than 0.2 dB/km loss, a figure of merit that means that dense wavelength-division multiplexing (DWDM) signals must be boosted by an optical fibre amplifier at least every 100 km. If they're not, then the signal-noise ratio becomes unacceptable.

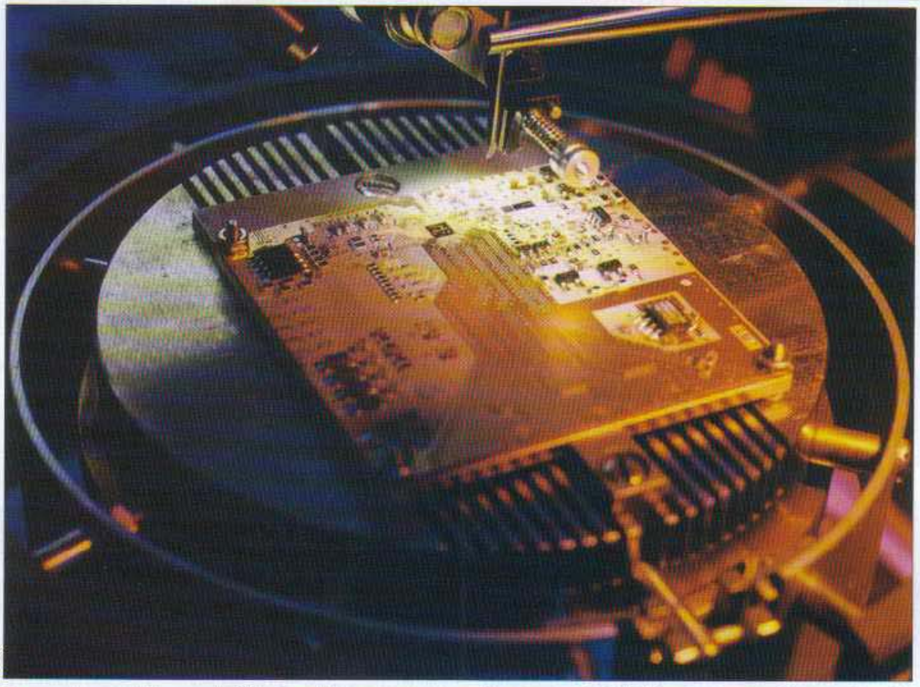

LUCENT

Converging worlds: most of the top-tier equipment vendors have 40 Gbit/s technologies (shown above) under development. The commercialization of such systems could pave the way for optical 3R.

Ultimately, however, other effects (see "Back to basics: the three Rs", p102) mean that the signal will be lost after a number of amplifier stages, depending on how cleverly the system is designed.

"Optical 3R is interesting once you have optical switching nodes," said Bernd Sartorius, senior research scientist at the Heinrich-Hertz-Institut (HHI) für Nachrichtentechnik in Berlin, Germany. "It makes no sense to do optical switching and then have optoelectronic conversion and electronic signal processing."

Industrial and academic research groups around the world have been working on all-optical 3R regeneration for some time. Although laboratory experiments have concatenated several optical regenerators, the approaches put forward are still highly sensitive to variations in input signal and the environment.

"Our lab tests use good-quality input signals and require extensive manual intervention to control and monitor transmission," admitted Sartorius. "We need to make sure that the scheme works for any signal and devise ways for the device to control its own operation."

Of course, that's where technologies like fibre Bragg gratings and vertical-cavity surface-emitting lasers were a few years ago. Now they're being sold by the crate-load. In each case, problems of product development were overcome one by one until the devices could be churned out affordably and in quantity.

However, optical 3R faces a significant hurdle before it can gain similar widespread acceptance. In short, how can the network operator be sure that a device operating at ultrahigh bit rates is working properly without including an expensive electrical monitoring system at that very bit rate – especially when the whole point of moving to all-optical regeneration is to avoid expensive high-speed electronics?

"When we first test [optical 3R] in the field, we will use expensive supervision, to make sure it works," confirmed Sartorius. "But it must have optical input, optical output and only low-cost, low-speed electronics. You need to work in the MHz, rather than GHz range, so there will have to be some sampling technique."

Sartorius and his team are well aware of the barriers to commercialization, but they're knocking them down one by one. Having demonstrated optical 3R in the lab, the HHI team supplied Alcatel with equipment and know-how for recirculating-loop experiments in which a 10 Gbit/s signal was transmitted over 15 000 km in the optical domain. This sort of collaboration is significant given that the final stages of optical 3R product development are likely to be taken up by a "big company", says Sartorius.

The price is right

Elsewhere, most top-tier telecoms equipment vendors now have 40 Gbit/s DWDM development programmes underway. And most, if not all, of them will offer 40 Gbit/s optoelectronic conversion in each network node. In parallel, there's also a concerted move by the vendors and the network operators to commercialize all-optical switching.

Such a convergence of events will, following Sartorius' logic, open the way for optical 3R regeneration. "People expect 40 Gbit/s to run with some expensive [electronic] parts," he said. "But if it's cost-effective, 40 Gbit/s is a good target to have optical 3R as a replacement for optoelectronic conversion."

Over in Japan, another research group working on optical 3R is looking further ahead. "Optical 3R will kick off at 160 Gbit/s per WDM channel," claimed Yoshiyasu Ueno, principal researcher in the Photonic and Wireless Devices Research Laboratories of the NEC Corporation. Certainly, all of the problems that make 40 Gbit/s electronics expensive will get a lot worse in 160 Gbit/s electronics. As well as the speed of operation, for example, Ueno says power consumption will become a limiting factor.

Optical signal degradation also happens faster and more severely as the bit rates climb. "Higher-rate signals tolerate relatively small timing-jitter accumulation, smaller polarization-mode dispersion and smaller group-velocity dispersion, all because of their relatively short bit separation," said Ueno. "Furthermore, higher-bit-rate signals suffer from larger cross-phase-modulation-induced noise."

It's worth noting, though, that the "speed limit" on the road ahead may not be 40 or 160 Gbit/s but, using some of the optical switching fabrics already on the market, a mixture of speeds from 10 Gbit/s upwards. This is another way in which all-optical solutions are preferable to their electronic equivalents – greater functionality.

"At one moment, the line will run at 10 Gbit/s and the next at 40, 80 or 160 Gbit/s," said Sartorius. "If you have fibre, you can operate at any speed and if you [optically] switch the signal, you will reach the concept of packet-switching." But that's another story (see page 107).

All-optical visions

So what will the promised land of all-optical signal processing look like? For one thing, true optical 3R will allow optical packets to travel around global networks indefinitely, until they find a route to their destination. In effect, that would mean a return to the early days of the Internet, when "bits" of e-mail traversed the network in an almost pseudo-random fashion until they reached their destination. Except, this time, the flow will be a multimedia flood, comprising massive data files, images, video and audio.

Another fledgling technology, known as Multiprotocol Lambda Switching (MPλS), should ensure that high-priority packets, like video- and audio-streaming, get through as quickly as possible. The number of "fat pipes" in the system should also ensure that very few packets take a significant detour *en route*, even during peaks in demand.

Whatever the details, optical 3R will certainly become a fixture in networks working at 80 Gbit/s per channel and higher, where the bandwidth demands of current applications could be catered for several times over. Projecting further, the technology is likely to trickle down to regional and then metro networks.

When this happens, a raft of new applications – video telephony, video-on-demand (assuming data storage keeps pace), remote gaming with ultrahigh-resolution graphics, subscriber access to sports and entertainment coverage from anywhere in the world – should all become both cheap for the subscriber and profitable for the service providers. In short, optical 3R will mean that the economic rules of global networking are likely to be rewritten once again – in much the same way that they were when DWDM made its entrance.

Back in the R&D laboratories, equipment vendors and telecoms operators are busy trying to figure out the optimum timescale for 40 Gbit/s deployment. But some researchers are already thinking generations beyond that. "After 160 Gbit/s, we will explore 640 Gbit/s," claimed Ueno. "Our technology is changing the viewpoint of optical materials researchers. In future, research on optically nonlinear materials will drive this technology up to and beyond 1 Tbit/s." ●

BACK TO BASICS: THE THREE RS

Signal regeneration, as the name suggests, is about rebuilding the data that was originally transmitted down a fibre-optic line. That means two things: an improvement in data quality and hence a reduction in bit-error rate.

OEO regenerators detect and recreate the original signal as new (figure 1). An all-optical regenerator could either do something similar, creating an altogether new data stream, or "tidy up" the incoming pulses. In the latter case, the process can more easily be broken down into three Rs: reamplification, reshaping and retiming.

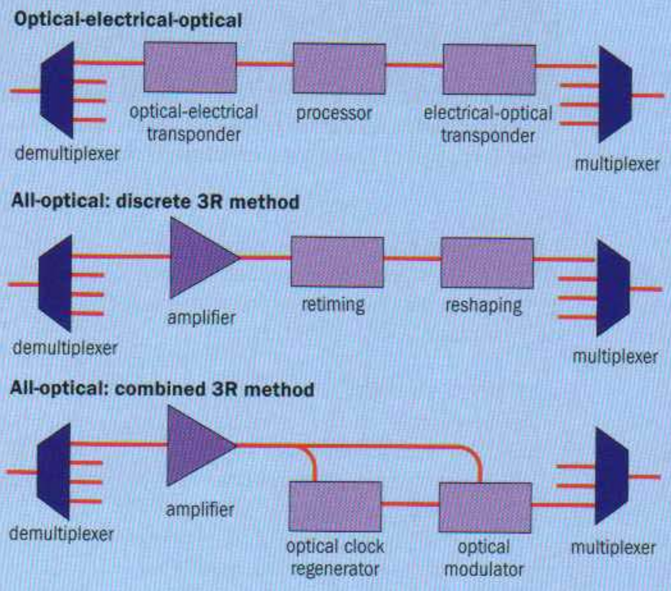

Fig. 1: Rival strategies for signal regeneration in long-haul networks.

Reamplification

The "Re" here is surplus to requirements. In principle, amplification does not improve signal quality, but rather recovers signals before they drop below detectable levels. Lasers are simply amplifiers with feedback, so for each type of laser there is also a type of amplifier.

Gain media, the signal-boosting materials, can be made of, or embedded in, semiconductor, glass, polymer, gas or crystal. In a compact optical 3R subsystem, the smart money is on semiconductor optical amplifiers to provide the amplification. Erbium-doped fibre amplifiers and Raman amplifiers are likely to remain as the workhorses of long-haul transmission.

The EDFA is currently the biggest-selling amplifier on the market. Erbium ions, impregnated in the core of a length of optical fibre, absorb pump light at 980 or 1480 nm and are stimulated by the passing signal to emit at around 1550 nm. The emissions are in phase with, and at the same wavelength as, the data signal, so the output is an amplified version of the input.

An EDFA can generate hundreds of milliwatts of output power (and over a watt is possible if the fibre is co-doped with ytterbium ions and enough pump power is available).

But interamplifier spans and the number of DWDM channels are limited by nonlinear effects rather than lack of gain. Scattered light, in spurious sidebands and broadband noise, is stimulated by high-intensity signals. The resultant noise can be amplified further down the line, so much so that it will eventually swamp the signal.

Another type of subsystem used to increase span lengths, and hence reduce the number of OEO regeneration sites in an optical link, is the Raman amplifier. A Raman amplifier uses one nonlinear effect (stimulated Raman scattering) to positive effect, as a booster for the signal within the transmission section of the fibre.

It's not a simple matter of boosting signals, though. Amplifiers also introduce photons to the signal at the wrong time and the wrong wavelength. Try as the amplifier designer might, there is an absolute limit to how noise-free an optical amplifier can be, set by the quantum-mechanical rules governing the behaviour of the gain medium. And the more gain, the more pain – in the form of amplified spontaneous emission.

The quality of the optical signals can be visualized by feeding them into a digital sampling oscilloscope and overlaying periodic portions on the screen. The result is termed an eye-diagram because, at the right setting, two round traces are seen next to each other on the screen.

As the signal quality deteriorates, however, the eyes close, making it harder to tell when a dot on the screen is associated with a "one" or a "zero", one timeslot or the next. This is where the other Rs come in handy.

Reshaping

Besides amplifier noise, optical pulses suffer from a host of other calumnies. Chromatic dispersion, polarization-mode dispersion and nonlinear effects all cause pulses to smear out.

Continued on page 104 >>

BACK TO BASICS: THE THREE RS

Since the cause of the problem usually has to be taken into consideration when looking for the solution, reshaping isn't as simple a proposition as reamplifying. There is one technique, however, that asks no questions when reshaping pulses and that is the use of a saturable absorber.

Saturable absorbers are materials that absorb low-level light more than higher levels. This usually involves quantum transitions in the material, which reset themselves in a shorter time period than the data clock-period. At low power levels, ions or molecules are always available to absorb passing photons, but at high power levels, every absorber is kept busy and the remaining light gets through.

Saturable absorption is a rather crude clean-up method, and not only from a purely aesthetic point of view. The leading edge of pulses tend to be eroded and steepened, while the trailing edge is left to drag out into a long tail. Conventional, non-return-to-zero (NRZ) pulses become lopsided and return-to-zero (RZ, of which solitons are a special case) become asymmetrical or get shifted out of their time slot.

The ideal reshaping subsystem would eliminate low-level noise, sharpen the leading and trailing edges, keep the top of NRZ pulses flat and RZ pulses bell-shaped. Unsurprisingly, no system does that perfectly – at least not yet.

Even then, just as the system designer has created beautifully shaped pulses, they start to wander out of their expected time window. That is to say, ones are detected where zeroes should be, and vice versa, because those ones arrive too early or too late. This is a task for the third R.

Retiming

The distinction between RZ and NRZ pulses is particularly important when it comes to timing because, although solitons have their benefits, they also have their drawbacks. Two of those drawbacks, soliton–soliton interaction and Gordon–Haus jitter, lead to timing problems.

The particle-like attributes that inspired the name "soliton" and attracted system designers in the first place are accompanied by other particle-like attributes that aren't helpful. For example, solitons of the same phase attract each other. The attraction pulls adjacent pulses out of their time slot and can ultimately lead to their destruction.

The other mistiming problem for solitons comes after amplification. Since noise is added randomly, the solitons are often lopsided by an infinitesimal amount after each amplifier. Such asymmetry disappears as the pulse propagates, in much the same way that a water droplet returns to its original raindrop shape after any perturbation.

Unfortunately, the new soliton centres itself on the new, average wavelength and proceeds at a slightly higher or lower speed. The researchers Hermann Haus and Jim Gordon calculated that the randomness of pulse arrival times, or "jitter", would grow logarithmically with distance.

Several cures have been proposed for Gordon–Haus jitter, all of varying practicality. One particularly ingenious idea is the use of progressively higher- or lower-frequency optical filters. Every soliton in the pulse train is then forced to recentre on the same optical frequency, marshalling them into the correct time slot and also removing noise. But there is no equivalent method to keep NRZ data and non-soliton RZ data synchronized.

Combined 3R

Rather than take the three Rs separately, many researchers, including Yoshiyasu Ueno of NEC and Bernd Sartorius of HHI, have tried to do the optical rebuilding job more directly – often in ways that are analogous to the OEO process. The aim of the game is to create a new signal without the need for high-speed electronics.

Instead of an optoelectronic clock-recovery circuit (used to synchronize the detector of the OEO system to the incoming signal), the optical–optical regenerator uses a method such as mode-locking of a distributed-feedback laser to create an optical bit-stream with precisely spaced pulses matched to the line rate of the incoming signal.

The second stage is to modulate the clock pulses with the original data and hence recreate the original signal. Such a direct approach has been taken by HHI. The output of its Q-switched laser has a sharp response to injected optical input, resulting in both reformed and reamplified output signals. The addition of modulation at the recovered clock frequency ensures that the output is also retimed. So far, the HHI experiments have yielded a negative power penalty of –2 dB on degraded pulses.

The interferometric alternative, being pioneered by Ueno's group at NEC, uses the data to gate pulses from a recovered optical clock stream. Such a precise interferometer is achieved by etching the waveguides into planar semiconductor. The width of the "gate" determines how far the pulses can stray from their proper time slot before they will be missed by the regenerator.

Whatever the technology used for optical 3R regeneration, solutions need to be developed from jury-rigged lab experiments into rugged subsystems that incorporate robust signal control.

FIBRESYSTEMS Europe

October 2001 Volume 5 Number 10

NEWS, TECHNOLOGY, APPLICATIONS AND PRODUCTS

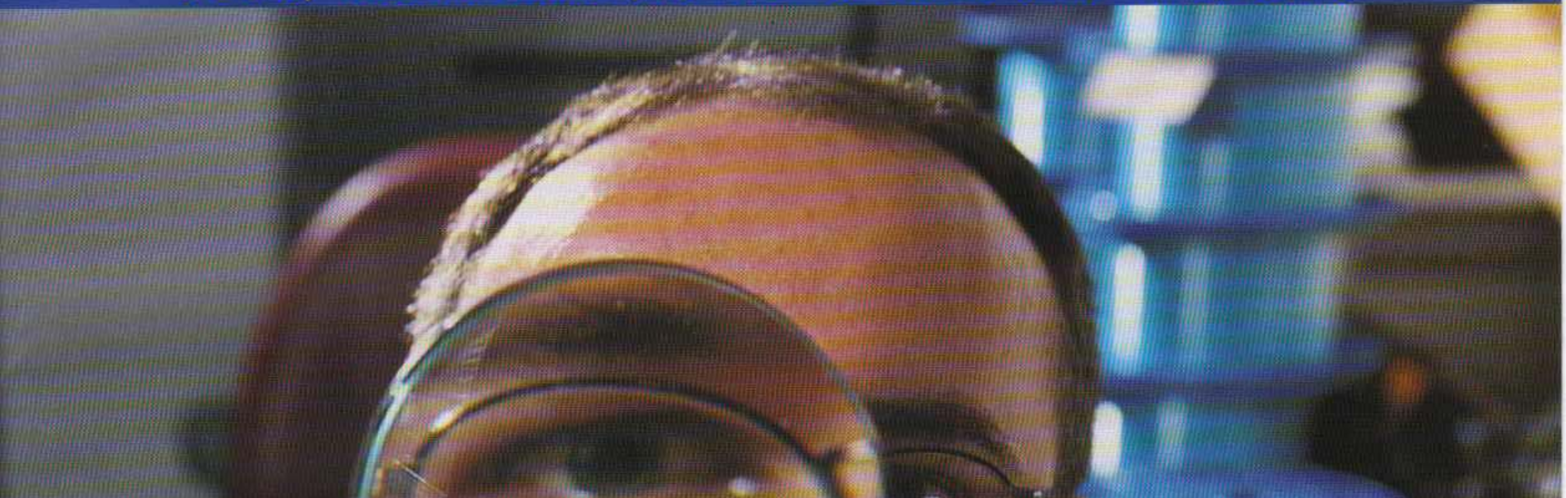

High-speed photonics: it's time to get dynamic

B2's broadband plan sure to hit home p55

ECOC 2001: special show preview p137

Opto vendors make automation pay p71