

LOW CROSSTALK DEVICES FOR WAVELENGTH-ROUTED NETWORKS.

M.J. Holmes* (now #), F.P. Payne*, P. Dainty#, T.J. Hall# and W.A. Crossland*.

Abstract: Over the last year, homodyne beat noise and bandwidth narrowing due to filter concatenation have been identified as major problems for large wavelength-routed networks. These problems will become more acute and will interact at the design stage, as wavelength channel spacings are decreased and as the channel bit rate is increased, leading to the requirement for devices with high fractional bandwidth per channel and with very low (<-50 dB) crosstalk. We explore the fundamental limits to crosstalk in optical routing components, and propose new design concepts for wavelength demultiplexers/multiplexers and space switches that have the potential to meet these strict performance requirements.

1 Introduction.

Wavelength division multiplexing (WDM) is an attractive technique for providing the high aggregate capacities into optical routing nodes because, with appropriate network design, the routing can be transparent to the bit-rate per channel and to the transport mechanism (PDB/SDH/ATM), and the routing does not require any form of synchronisation between channels [1]. In the next section we outline two of the current major problems in wavelength-routed networks, and describe how these problems lead to the demand for wavelength-routing devices with high fractional bandwidth per channel and very low crosstalk, typically - 50 dB. The purpose of the work presented in this paper was to investigate the design of wavelength-routing components with such properties.

It is an aim of the POETS project to investigate free-space implementations of optical routing, based on the use of fixed and dynamic holographic components. In the context of a colloquium on guided-wave devices, we were interested in answering the following questions: what sort of crosstalk (and bandwidth) performance might we be able to achieve with holographic (free-space) optical components, how does the likely crosstalk performance of the holographic devices compare with that of guided-wave devices, and how can we take advantage of recent advances in microengineering techniques [2], and exploit the potential for 2-D fan-out with free-space optics, in order to design compact devices.

In section 3 we discuss how the various choices made during the design of a space-switch will influence the final crosstalk, and compare the fundamental limits to the crosstalk in guided-wave and free-space optical switching: we predict that for semiconductor integrated guided-wave devices there is a 'background' level of crosstalk, induced by scattering from surface roughness,

even for switches with infinite extinction ratio. For 'free-space' optical switching we show that the 'background' crosstalk may be suppressed by exploiting the coupling behaviour of Gaussian beams into single-mode fibres.

In section 4 we present a new holographic implementation for optical switching with spatial light modulators, that has been designed to eliminate the crosstalk from higher diffraction orders. In section 5 we show that guided-wave wavelength demultiplexing and remultiplexing devices are unlikely to provide the combination of high fractional bandwidth and low crosstalk that will be required in wavelength-routed networks. We also discuss the problems that arise for a blazed-grating wavelength demultiplexer when narrow channel spacings are required. Finally, in section 6 we present a new design of wavelength demultiplexer, designed to overcome the problems of a blazed grating approach.

2 Current problems in wavelength-routed networks.

One of the first applications of optical routing is likely to be an evolutionary one: not a fully transparent optical network but a transparent optical transport layer, overlaid on the electronic transport layer [3,4]. The function of the optical crossconnects would be to route high capacity tandem traffic, bypassing the electronics, and 'adding (dropping)' the lower capacity channels and the originating (terminating) traffic from (to) an electronic cross-connect. The function of a WDM crossconnect is to set up semi-permanent routes for each wavelength channel: this may be achieved with three optical stages: a wavelength demultiplexer on every input fibre, followed by a reconfigurable space-switch, followed by a wavelength multiplexer on every output fibre (figure 1).


As a result of recent demonstrator project

* Cambridge University Engineering Dept., Trumpington Street, Cambridge CB2 1PZ, England.

Department of Electronic and Electrical Engineering, King's College, The Strand, London WC2R 2LS.

wavelength-routing nodes [5], two major problems have been discovered to afflict wavelength-routing networks. The problems are bandwidth narrowing and homodyne beat noise, and they interact at the design stage. Wavelength routing necessarily involves a wavelength selection process (in the demux stage and perhaps the mux stage): every wavelength selection process has a finite filter bandwidth, and each subsequent filtering operation leads to a narrowing of the net bandwidth perceived by a routed channel. For a large wavelength-routed network the net end-to-end bandwidth can become very narrow, leading to the demand for tight control on the transmitter wavelengths: good wavelength stability can be achieved with fibre gratings [6]. In a large network we would then have many transmitter lasers at almost (within the wavelength referencing tolerance) the same wavelength. As a signal traverses the network, crosstalk in the routing optics will lead to the accumulation of in-band crosstalk, originating from other transmitters at the same (nominal) system wavelength. Because the receiver is a square-law device, the crosstalk will beat with the signal. If the frequency difference between the signal and crosstalk is within the receiver bandwidth, the crosstalk will corrupt the data.

Figure 1: Transparent Optical Cross-Connect



Even without tight wavelength referencing, homodyne beat noise will still arise due to crosstalk originating from the same transmitter as the signal itself: this occurs due to non-perfect wavelength demultiplexing and remultiplexing [7], and depends on the routing configuration.

The net crosstalk amplitude can be reduced by making the multiplexer wavelength dependent [8], but this requires an extra filtering operation per routing node, and so will exacerbate the bandwidth narrowing problem. Reducing the device crosstalk will very often lead to a reduction in the allowed filter bandwidth, again making the bandwidth narrowing worse. A narrower end-to-end bandwidth requires tighter wavelength referencing, leading to worse homodyne beat noise.

The conclusion from this cycle, is that for the wavelength filtering devices we should seek to maximise the bandwidth we can achieve for a given crosstalk, and for the space-switches we should seek to minimise the crosstalk.

We have used a statistical approach to simulate the accumulation of beat-noise terms for an optical transport network with as many nodes as the current UK inner-core network, and have found that (for a maximum receiver penalty of 2 dB at error rates of 1 in 10^9), the crosstalk requirements vary between -43 dB and -50 dB, depending on the network architecture.

3 Space switch design for low crosstalk.

For wavelength-routing the space-switches need reasonable fan-out, e.g. 4 or 8, rather than fast switching speed: for example the reconfiguration time for an electronic (SDH) crossconnect is 20 ms [9]. The routing configuration is controlled by electronic signals sent by the local element management centre: all-optical switching is not necessary. Therefore very fast nonlinear 'all-optical' switches are outside the scope of this study, although they may have other roles to play in WDM networks for use as wavelength converters, for example.

The design of a space-switching component breaks down into 3 stages: the first is the choice of a process or method for performing a space-switching function; the second is the choice of a particular architecture or arrangement of the sub-switch components; the third stage is the choice of a particular implementation: that is the device technology and the details of the device design. Decisions made at all 3 stages of the design process have implications for the final crosstalk levels.

3.1 Switch method.

(i) **Guided-wave switches:** Guided-wave switches fall into two classes: those based on interferometers, and those performing a more digital, 'gating' function. The interferometer-based switches are operated by adjusting the effective indices of parallel waveguides. This process will inevitably be prone to high crosstalk, because small changes in the effective index of one guide can lead to large changes in the power coupled across. Possibly the best crosstalk results that have been obtained with an interferometer method are -24 dB crosstalk in a 4 by 4 matrix switch using electro-optic effects in a directional coupler [10].

The 'gating' class of guided-wave switches is less crosstalk-prone: one method involves splitting the input power so as to take several 'copies' of the input signal. The passage of one copy of the signal in a particular direction is controlled by turning gain blocks on and off: in the 'on' state the gain compensates for the splitting loss, and in the 'off' state the signal is blocked by the attenuation of the gain block. Both semiconductor laser amplifiers and rare-earth doped fibre amplifiers have been used as gain blocks to perform

a switching function. Another 'gating' switch is based on the 'digital optical' Y switch. The crosstalk for this class of switch depends on the extinction ratio of each individual switch element. For 'gating' switches implemented into integrated semiconductor devices, reported extinction ratios for the SLA are 40 dB [11], and for the Y-switch are 40 dB [12].

(ii) Free-space switches: 'Shadow-routing' switches are the free-space equivalent of the semiconductor laser amplifier guided-wave switch, except that the shadow-routing uses attenuation instead of gain to control the routing. For the shadow-routing, a fixed hologram acts as a splitter to take many copies of the input signal. All but the chosen copy are blocked with an (amplitude) spatial light modulator: at present typical modulator contrast ratios are in the range 150 to 200, but experimental studies indicate that an order of magnitude increase in the contrast ratio is possible with refined substrate properties. The shadow routing switches can have very high fan-out: for example a 64 by 64 crossbar switch has been demonstrated as part of the OCPM project [13]. However, to achieve such high fan-out requires the use of multimode fibre in the output plane, and this would preclude the use of such switches in a transparent optical network.

The second free-space switching method for performing a switching function is to 'beam-steer' the input signal to the required output waveguide. This method is perhaps the least prone to crosstalk out of all four (guided-wave+free-space) methods discussed in this paper: assuming we have steered the beam to the correct output, the resulting crosstalk will come from the evanescent tails of beams steered to an adjacent output port. For Gaussian output beams of spot size 'x', matched into output waveguides spaced distance 's' apart, the theoretical crosstalk is $-4.34 (s/x)^2$ (using [14]). For 'standard' type telecomms fibres spaced 250 μm apart, and planar silica waveguides spaced 50 μm apart, this theoretical crosstalk is only -10,400 dB and -940 dB, respectively (!). This is far too small to be measured and would not contribute to homodyne beat noise problems in even the largest of networks. Examples of beam-steering switches include 'Star' fibres, acousto-optic beam deflectors and liquid-crystal holograms.

3.2 Switch architecture.

It is well-known that dilated switch architectures will reduce the net crosstalk for a space-switch. For example, a logical N by N (crossbar) switch can be formed from a 1:N switch at every input port, fibre 'wired' to N:1 switches at every output port. For this architecture, 2 'crosstalk events' must occur for crosstalk to appear at the output. Hence the net crosstalk is second-order:

a net crosstalk of -60 dB for a space-switching stage can be achieved with 1:N and N:1 switches with crosstalk between ports of -30 dB.

A secondary advantage of a dilated architecture is that it is easily upgraded: a dilated N by N switch can be progressively upgraded to a MN by MN switch by placing 1:N switches at each of the o/p's of the 1:N switch.

3.3 Switch implementation.

The choice of device technology can also influence the crosstalk.

(i) Guided wave switches: The lowest crosstalk guided-wave switches are those using semiconductor 'gating' elements. These are integrated into a planar semiconductor device, with connection between the gain/loss blocks in waveguide form. The semiconductor device fabrication process can introduce imperfections in the waveguide walls. The typical feature size for these imperfections is close to the carrier wavelength for optical signals. We were interested to see whether scattering from these surfaces would lead to a significant 'background' level to the crosstalk, that would occur even for switch elements with infinite extinction ratio.

In this technology, significant attenuation in the waveguides occurs as a result of mode coupling from the fundamental waveguide mode(s) to radiation and substrate modes, where the mode coupling is excited by the surface roughness of the waveguide walls. Typical loss coefficients for this scattering mechanism are between 3 cm^{-1} and 5 cm^{-1} . The coupling can be interpreted as being equivalent to a given probability (per unit length) of a photon being coupled out of the waveguide. On reaching an adjacent waveguide, the photon will have the same probability of being coupled into this waveguide. Hence the mode coupling will lead to a 'background' level of crosstalk, even for switch blocks with infinite extinction ratios. Earlier theory developed to calculate the attenuation due to this surface roughness [15], has been adapted to calculate the crosstalk. It was found that the ratio of the (absolute) crosstalk, C, to the square of the loss coefficient due to this scattering mechanism is given by:


$$C / \alpha^2 \propto L^2 / D \quad (1)$$

where L is the length of parallel waveguides, D their separation, and the constant of proportionality depends strongly on the correlation length of the surface roughness, but is insensitive to all other parameters. The maximum crosstalk occurs at short correlation lengths of the order of 0.05 μm : with a loss coefficient of 5 cm^{-1} , and a 1 μm length of parallel waveguides separated by 250 μm , we estimate the crosstalk to be -50 dB.

(ii) Free-space switches: Reflections in a free-space optical system can lead to two distinct crosstalk mechanisms: reflection into the wrong o/p port (not the intended o/p), leading to a crosstalk of say, -C dB per switch, can be reduced to a net crosstalk level of -2C dB, given a dilated switch architecture. However, coherent (unwanted) reflections into the intended o/p port, cannot be removed with a dilated architecture. The coherent scattering should therefore be less than -50 dB, while the adjacent channel crosstalk can be -25 dB with a dilated architecture. Unwanted diffraction orders in a holographic system can also lead to crosstalk.

In a free-space holographic optical system, unwanted reflections will occur from the input and output lenses, the hologram surfaces, the cleaved fibre ends, and the fibre mounts. High-quality commercial AR coatings on the lens surfaces will bring these reflectances down to less than 0.12% (equivalent to -29 dB) over a 40 nm window [16]. For Fresnel reflection off the fibre end we assume an effective fibre index of 1.445, giving a net reflectance of 3.3% or -14.8 dB. One side of the hologram will be rough, due to the devices/pixels used to form the hologram: it would be difficult to AR coat this surface so we assume a reflectance of around 4% (-14 dB). The other side of the hologram could be AR coated with a reflectance of -29 dB (as for the lens surfaces).

Figure 2: Hologram used in transmission.



For a holographic system used in transmission (figure 2), two reflections must occur for coherent scattering or crosstalk to be coupled into the o/p fibres, so that the net effect is a second-order function of the reflection from a single surface. In order to maximise transmitted power and minimise spherical aberrations we assume the use of precision-moulded plano-aspheric lenses for the i/p lenses, with the planar surfaces closest to the fibre ends. We have estimated the size of the reflections from each possible pair of surfaces. Single-mode fibres will only accept light from i/p beams that are well-focused, at near-normal incidence, and with a beam centre close to the fibre core: we have also estimated how much of the reflected power is coupled into the input and output fibres, using the standard formulae [14] for the coupling efficiency of Gaussian beams into standard telecoms fibres.


Beams reflecting from a lens surface will diffract and be defocused on the fibre ends: the resulting phase-front

curvature leads to very weak coupling into the fibre, therefore most reflection pairs including a reflection from a lens surface will cause negligible coherent scattering. The strongest event in this set occurs for light that is initially reflected from the cleaved end of the output fibre, then again reflected from the near surface of the output lens. After reflection from a plane lens surface and coupling back into the o/p fibre, the net crosstalk level is estimated to be -80 dB, for a lens surface 5 mm from the fibre end.

Beams reflecting from the rough hologram surface: normal reflections from the hologram surfaces will be refocused by the lenses and can therefore be strongly coupled into the i/p and o/p fibres. We assume that the rough surface of the hologram is facing towards the i/p fibre. Reflection from this hologram surface, followed by transmission through the i/p lens, will lead to a significant back-reflection into the i/p fibre. In a dilated switch architecture, two such scattering events will lead to coherent scattering levels of -28 dB or more, depending on the roughness of the hologram surface, leading to significant beat-noise. Reflection from the rough hologram surface, followed by reflection from the cleaved end of the input fibre would also lead to -28 dB coherent scattering in the output fibre. However, we can avoid these crosstalk mechanisms by placing the i/p fibre slightly off-axis (figure 3). For an i/p fibre offset by 'o' μm, the reflected beam would be offset by 2'o' μm, and the fraction of backscattered power coupled into the i/p fibre would be $-4.343(2o/x)^2$ dB [14]: an offset of 11 μm is sufficient to suppress the backscattered power by 80 dB.

Figure 3:

Reflected beam paths with angle-polished fibre ends and i/p beam in off-normal incidence to the hologram.



Beams reflecting from the cleaved end of both fibres: the estimated crosstalk for this mechanism would be -50 dB. This effect can be suppressed by polishing the face of the input fibre: an angle of 8 degrees is known to give the best compromise between (out)coupling loss and suppression of backscatter.

Reflection off the cleaved end of the o/p fibre, followed by reflection from the smooth (and AR coated) surface of the hologram (figure 3) may lead to crosstalk into other o/p fibres, depending on their position. With worst-case positioning, the crosstalk would be -43 dB.

With careful choice of the position of the o/p fibres, this crosstalk could be reduced considerably. A 15 μm separation between the centre of the crosstalk beam, and the nearest o/p fibre, would reduce the crosstalk to below - 80 dB. Alternatively we could polish the end of the output fibres to suppress the reflection.

Crosstalk from unwanted diffraction orders.

The crosstalk and coherent scattering due to reflections will occur in any implementation of a beam-steering hologram. The crosstalk from unwanted diffraction orders depends on the specific details of the hologram technology and design.


Holographic beam-steering can be implemented with phase modulation of a spatial light modulator (SLM) [17]. Polarisation-independent operation [18] can be achieved with a binary phase hologram, formed from a 2-D pixellated array of ferroelectric liquid-crystal, embedded in, and controlled by, a VLSI silicon backplane consisting of 2 μm CMOS. Binary-phase holograms are so-called because they can induce two different values of path difference in light passing through the pixels. The relative phases are usually 0 and π , and are controlled by rotating the liquid-crystal molecules, so as to adjust the refractive index experienced by light passing through the liquid crystal. The fraction of incident power diffracted by the hologram varies as $\sin^2(2t)$, where t is half the angle through which the molecule is rotated. Half-angles of 36 degrees have recently been achieved [19], with a switching time of 80 ns: such devices will allow diffraction of 90 % of the power incident on the pixels. These devices are an attractive component for future telecomms networks because they are potentially very cheap, they require only standard 10V digital supply voltages, and the 2-D operation allows a large fan-out per switch. Other examples of beam-steering switches are acousto-optic beam deflectors, which require RF supply, and 'Start' switches, which are limited to 2:2 (crossbar) operation, although bigger switches can be made by cascading many 2 by 2 crossbars.

An SLM is used as a beam-steerer by changing (electronically) the phase of chosen pixels in order to construct a phase diffraction grating with a tuneable period and pattern. For pure beam-steering we require a perfect sawtooth phase diffraction grating: for this case we would get diffraction into a single grating order, and the output angle of the switched light would then be given by:

$$\sin \theta = \lambda / Q \quad (2)$$

where Q is the grating (sawtooth) period. By changing the sawtooth period we change the output angle, and switch the output between different waveguides (fig 4).


Figure 4: Principle of holographic beam steering



For binary-phase operation we are limited to 2 discrete phase levels, so we cannot form a sawtooth phase variation. The closest binary-phase approximation to a sawtooth is a 'square wave', as shown in figure 5, with equal width stripes inducing alternate phase shifts of 0 and π . For this case (and with 1-D fanout) 80% of the input power is diffracted into two (equal) main orders, positioned symmetrically about the optical axis. The rest of the power goes into higher-order grating modes: the relative amplitude of each mode is shown in figure 6, where the output angle of the m 'th grating order is given by:

$$\sin \theta = m \lambda / Q \quad (3)$$

Figure 5:



For a square-wave SLM, we choose the grating period such that the light diffracted into one of the first-order grating modes is coupled into the selected output fibre. Unfortunately, the higher-order grating modes will then lead to severe crosstalk whenever the light diffracted into these orders is coupled into another (NOT selected) output waveguide. One method to suppress this crosstalk is to change the grating pattern: from a square wave to a more complex (computer-optimised) structure, designed to suppress the higher-orders. Crosstalk levels of -35 dB have been achieved by this method, but the penalty is that a large number of pixels are required in each period, leading to a small output angle for a fixed pixel pitch, and consequently long devices.

Acousto-optic beam deflectors can also be used to implement a free-space optical switch. We have not calculated the crosstalk levels in these devices but note that the acoustic wave is usually at around 100 MHz. Hence at room temperature we would expect 50,000

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.