Optical Data Communication

Technology, Standards, Markets and Future

The technology of optical data communication is summarised in the fields of optical fibre, optical sources and detectors, multiplexing technologies, assembly techniques and mechanical and electronic technologies.

The optical-data-communication marketplace is driven by interconnection standards and multi-source agreements between suppliers. Once products are established, competing with them is difficult except on price. The opportunities for new entrants are with new products while the new standards are developed; the standard-setting process is a marketing opportunity in itself, giving participants some influence over whether standards include their technologies as well as a chance to meet customers and assess competitors.

The marketplace is international and is volume driven, but it is not exclusively dominated by large corporations; many SMEs are active and examples are given of typical market opportunities.

There are numerous future opportunities in this marketplace. The examples of fibre to the home and microprocessor interconnections are described.

Judith Hankey

Jurion Technology

PRIME Faraday Partnership

Published in 2004 by
PRIME Faraday Partnership
Wolfson School of Mechanical and Manufacturing Engineering
Loughborough University, Loughborough, Leics LE11 3TU
http://www.primetechnologywatch.org.uk/

ISBN 1-84402-030-4

Comments on this publication are welcomed. Please send them to
<feedback@primetechnologywatch.org.uk>

Contents

1.0 Introduction

One transatlantic fibre-optic cable has an enormous information-carrying capacity. Each of its many fibres can handle in the region of 100 gigabits per second (Gb/s) – enough to carry about 10 billion characters of text per second or 10,000 video channels or 1 million telephone calls. The total capacity of the internet backbone is similar in size to the telephone network capacity of the USA and totals roughly 2 terabits per second.[1] Transatlantic capacity well exceeds the present USA traffic capacity and these long-haul telecommunications systems are presently under-utilised. For each high-capacity long-distance fibre-optic link, there are significantly more slower shorter links which would ultimately feed the longer-distance network if the traffic was international. Each PC or other data terminal must have the connection in place even if is not used at that moment, so the total market for links is much larger than the backbone capacity. A significant proportion of traffic is local, between PCs and the server, in the same building or campus, and will never transit the countries' backbone, so the largest volume market is in the lower data rate links that are widely geographically distributed.

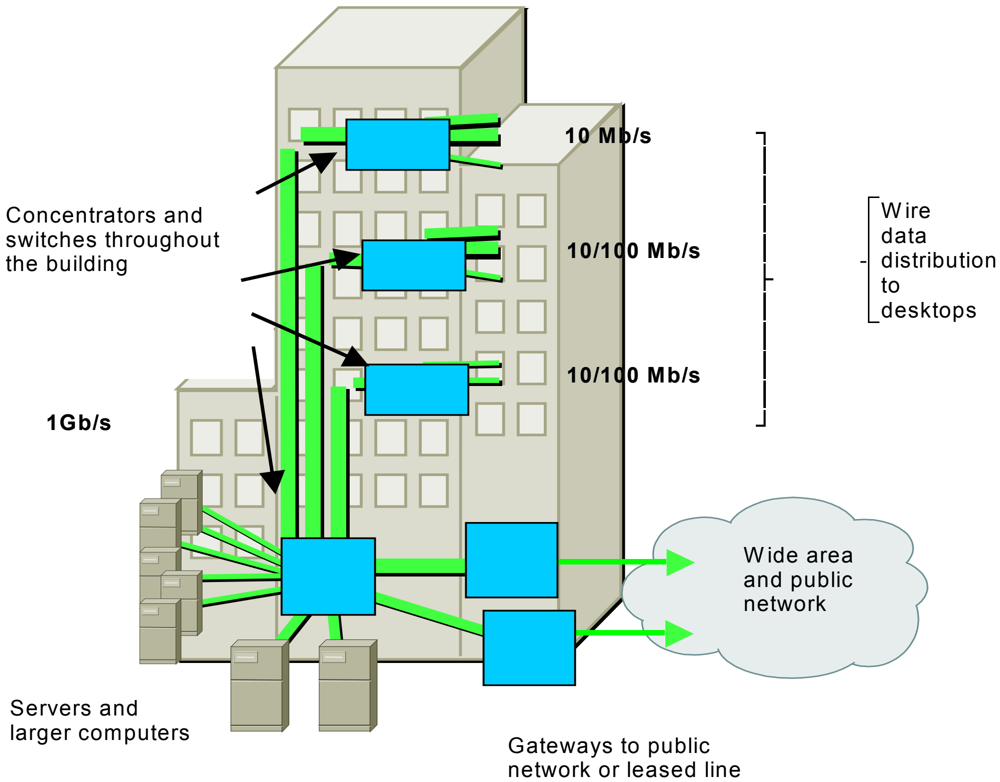

Figure 1 Data communications within a typical large building. Links carrying data above 100 Mb/s are likely to be optical.

Although virtually invisible to the end user, optical technologies are widely used in data communications. Take the typical large office with a PC on every desk. Those PCs are normally connected to its network by a copper cable. This wiring leads to a concentrator, from which optical interconnections are frequently used between floors of a building and, on the building backbone, to the typical computer room (Figure 1). The dominant networking technology used for these connections is Ethernet and optical communications fit into this standard. The sheer volume of networked PCs across the globe dictates the size of the market for networking components and optical data-communications components are one sector of that market. This report explains the fundamentals of optical data communications and describes in further detail what is involved and how the technology is applied.

Figure 2 shows the schematic of a simplified (uni-directional) optical link. (In reality links are not sold as a unit but as transmitters or receivers, used on separately installed fibre, or as transceivers, single modules containing both a transmitter and a receiver and allowing two-way (duplex) data transmission.)

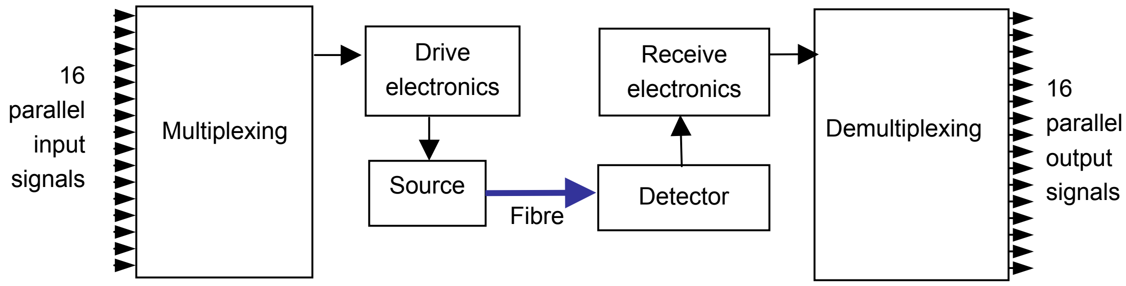

Figure 2 Block diagram schematic of an optical link

The digital data transmitted over optical fibres, shown in the middle of the figure, may arise from voice traffic or audio and video files. The specialist issues that arise when the information is translated from analogue form or the equally specialist telecommunications techniques are not covered.

The next chapter of this report outlines this fundamental technology in more detail, starting with optical fibres, since they determine all the other technologies used. This is intended as an overview of the concepts so that the reader can follow later sections or research in more depth by referring to the texts listed in the further-reading section (p. 36).[2] The mechanical technology used in data communications is covered in Chapter 3 and the electronics in Chapter 4. How the mechanical and electronics are interfaced together and the trade-offs between performance and price for a complete transceiver module are described in Chapter 5.

The nature and present state of the market are determined by standards which ensure that each optical link will always work even when different companies from all over the world manufacture each component. As a consequence, this is a commodity-style market place, and components are bought primarily on cost. The resulting competition drives down prices and ensures that sales volume and penetration continue to increase. Fibre connections to the desk and fibre to the home are not commonplace now but as data rates increase and costs come down then, they will become commonplace. The market is described in more detail in Chapter 6 together with the dominant current and emerging standards in Chapter 7. The future technical opportunities in the field and examples of SMEs in the present supply chain are given in Chapters 8 and 9 respectively.

2.0 Fundamentals of Optical Technology

2.1 Fibre

Optical fibre has become the transmission medium of choice because it has a very low loss as the signal is transmitted along it, typically 0.2 dB per km, or 5% for every kilometre of travel. Thus after 15km of fibre half of the light power is still available to be received.

2.1.1 Loss

Generally optical loss is reduced for longer wavelengths, and as component technologies have matured optical, sources and detectors have become available at longer wavelengths, reducing losses, and lengthening the distances that can covered in a single span. Early technology allowed sources with wavelengths of 800–900 nm based on the semiconductor gallium arsenide (GaAs) and later at 1300 nm and 1550 nm based on the semiconductor indium phosphide (InP). Wavelengths between the 1300-nm and 1550-nm windows have become accessible as improvements in fibre purity have reduced optical absorption due to the water content of the fibre. See Figure 3. A useful loss comparison can be made with copper coaxial electrical cable (of higher quality than the coaxial cable that typically connects a TV set to its antenna). For coaxial cable, the distance over which the electrical power is halved is 1 m; for optical fibre, it is 15 km. Another reason why fibre has become the medium of choice is cost: bare fibre costs only 10 cents (US) per metre in bulk; good-quality coaxial cables cost $10 per metre.

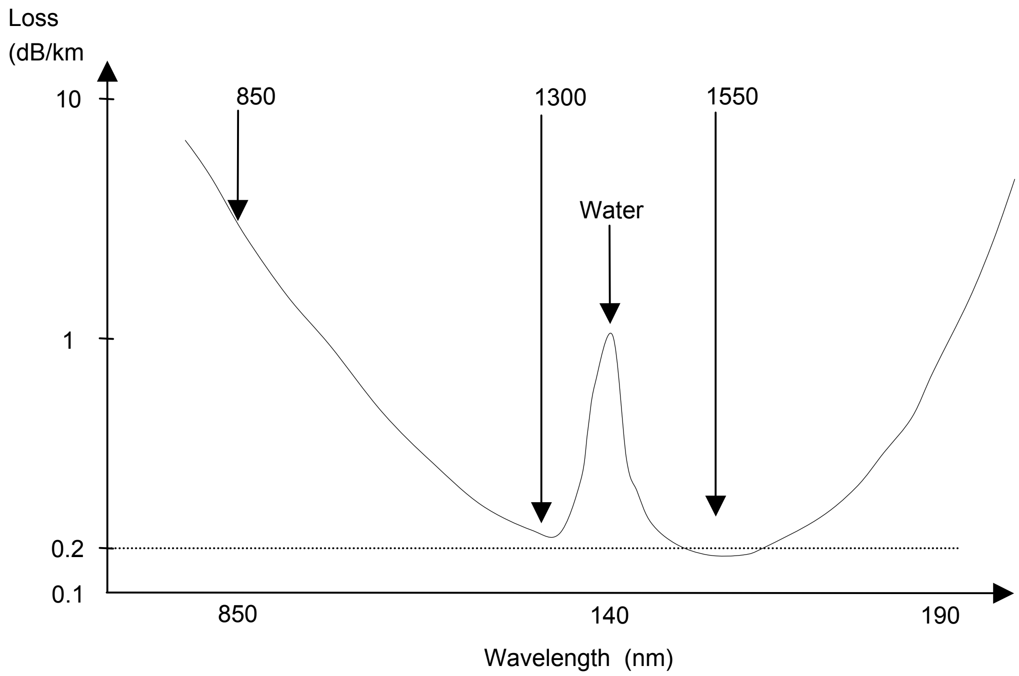

Figure 3 The variation of optical power loss with the wavelength of the light used

2.1.2 FIBRE STRUCTURE

The structure of an optical fibre determines its performance. Optical fibre is composed of two types of glass (Figure 4). One is used for the central core, in which the signal travels, the other for the cladding, which keeps the signal inside the core. The optical property that keeps the signal in the core is total internal reflection (the phenomenon responsible for your not being able to see below the surface of a lake, for example, when your line of sight is less than about 5° from the horizontal). This is achieved by doping the glass with different elements and creates a difference in optical refractive index, normally denoted by the letter ‘n’. The glass is fabricated in a cylinder and then drawn whilst molten (‘pulled’) to an outer diameter of typically 0.125 mm. The fibre is then protected by several polymer jacket layers to allow normal handling, and if required, formed into a cable with several additional fibres.

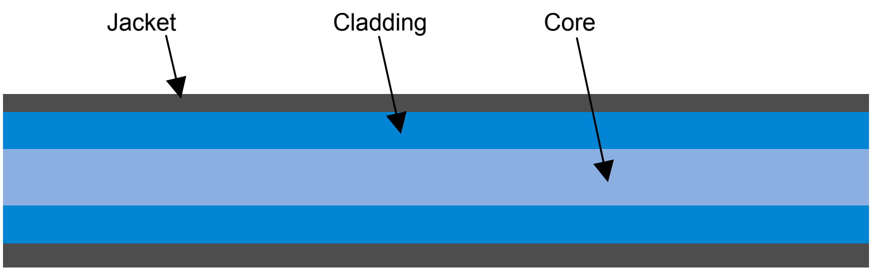

Figure 4 The structure of an optical fibre

There are two main types of fibre: single-mode, with a core thickness of approximately 7μm; and multimode, with a core thickness of 50μm or larger (Figure 5). (There have been several different standards historically, up to 200 μm, but the newer standards for future use are 50 μm, since it has higher bandwidth).

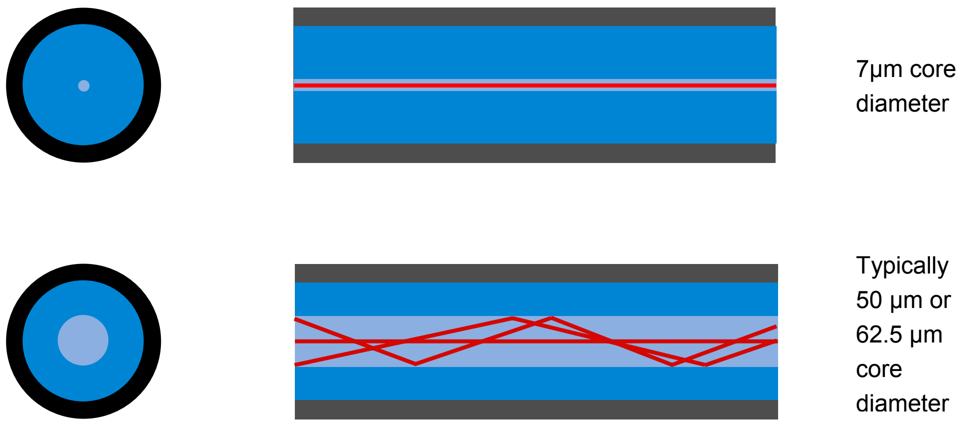

Figure 5 A single-mode and a multimode fibre (with light paths shown in red)

The advantage of the smaller core size, which has more demanding assembly tolerances, is the higher bandwidth for longer-distance transmission or higher data rate. Single-mode fibre has a higher bandwidth than multimode because the light can travel only in a single pathway (or ‘mode’) within the core. Multimode fibre can support several tens to hundreds of different modes in the core, which means each light beam can travel a different distance depending on whether it goes straight through the centre or it bounces from side to side between the input and output ends.

If a light pulse is injected into the fibre, since light speed is the same in each mode but path lengths differ between modes, parts of the pulse have to travel a longer distance than others and reach the output end later. A sharp square pulse going into the fibre becomes ‘smeared out’ in time and can spread into the next pulse and stop the individual bits of data from being distinguished (Figure 6). This effect is known as ‘modal dispersion’ and gives rise to an upper limit on the rate at which data can be transmitted, because the distinctness of adjacent pulses becomes increasingly obscured by the smearing out as frequency increases.

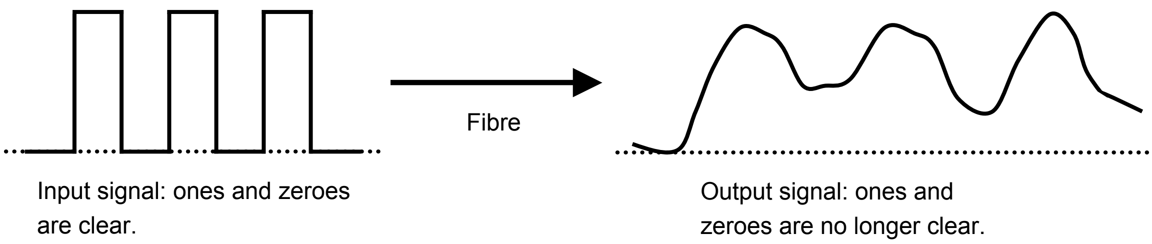

Figure 6 The distortion of data pulses due to modal dispersion in multimode fibres.

In addition to modal dispersion there is another limiting effect, known as ‘chromatic dispersion’, due to differences in the speed at which different wavelengths of light travel through the glass of the core (Figure 7). It affects single-mode and multimode fibres.

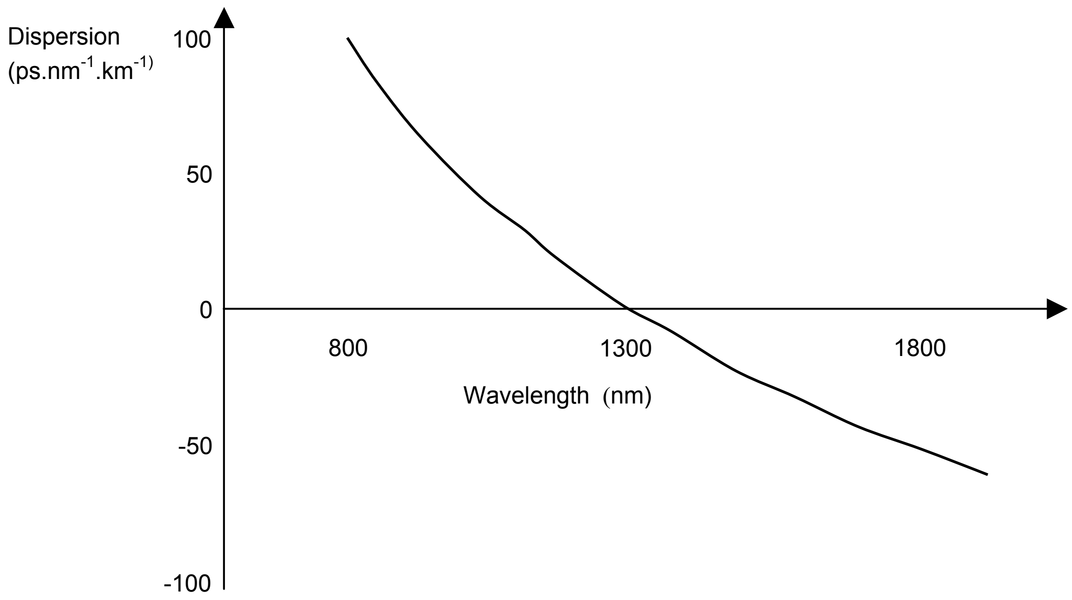

Figure 7 The variation of chromatic dispersion with wavelength

For most types of fibre, chromatic dispersion is zero at around 1300 nm. The highest transmission capability (or bandwidth) achievable is with single-mode fibre at this wavelength. At 1500 nm, loss is lower but chromatic dispersion is greater; long-distance or high-frequency applications at 1500 nm thus require dispersion compensation.

2.1.3 RIBBON

Ribbon fibre is an array of normal multimode fibres accurately placed alongside each other to make up a ribbon similar in outline to copper-wire ribbon cable. Ribbon fibre is typically 50µm multimode type glass, and is mostly used at 850 nm. But in future 1300 nm may be viable. As with copper ribbons, the main advantages are ease of mechanical handling and removal of the need for cable ties, coupled with the use of a single multi-way connector which will be lower cost and less bulky than many single-way connectors. The use of array connectors[3] ensures the connector assembly is only marginally more expensive than single-way styles.

2.1.4 PLASTIC

Plastic fibre has been available for many years and has the advantage of cost and ease of connection, but the optical attenuation (loss per metre) performance is not as good as for glass; PMMA-type plastic fibre has an attenuation of 150 dB/km at 650 nm, for example. So it is not suitable for use over longer distances. However there has been a lot of progress,[4] which may allow more extensive deployment of plastic fibre over shorter distances, either within the home or to directly connect integrated circuits. There has been considerable work recently in the automotive industry investigating the use of plastic fibre[5]–to transmit information for in-car entertainment or from sensors, to carry control data around the car and even to carry light from a single source to multiple external lights.

2.2 Sources

An optical source consists of a diode that emits light when an electric current is passed through it. Laser sources are the most common for present designs since they give a higher output power than the simpler light-emitting diode of previous products. This section describes the basic types. There are many variants.

2.2.1 SURFACE EMITTING LASERS FOR MULTIMODE FIBRE

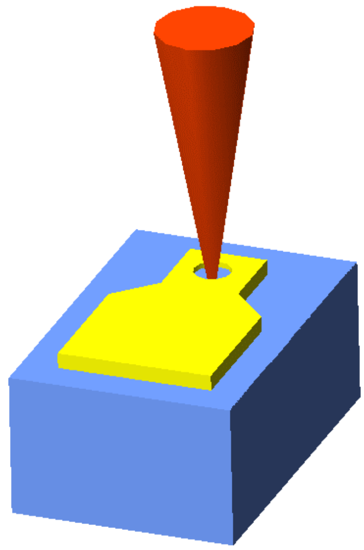

Figure 8 Schematic of a surface emitting laser

Figure 8 shows a schematic of a vertical-cavity surface-emitting laser (VCSEL). The light source is formed in the semiconductor wafer and emits up out of the surface. The advantage of this structure is that the devices can be tested on the wafer before they are cleaved into individual chips and this assists in keeping the cost of the chips low. The wavelength of the light emitted has for several years been limited to 850 nm but at the time of writing (mid-2004) manufacturers expect soon to be producing 1300-nm VCSELs of sufficient output power and reliability to compete with edge-emitting lasers (see below) at 1300 nm. The shape and size of the light beam from VCSELs is very well suited for multimode fibre which is the most common installed fibre in buildings. These sources are one of the building blocks of optical data communications and have the advantage that they can easily be built as arrays for parallel fibre coupling into ribbon cable. They have low electrical power consumption and medium optical output powers. However it is not the output power which limits the length the signal can travel but the inherent dispersion at the operating wavelength of 850 nm. This length limitation depends on the data rate but can be around 300 m at 1 Gb/s.

2.2.2 EDGE EMITTING LASERS FOR SINGLE-MODE FIBRE

Figure 9 shows an edge-emitting laser, which is fabricated on a semiconductor wafer. The device is neither complete nor testable until the chip is cleaved from its wafer. The chip is usually approximately 0.5 mm square and 0.1 mm thick and needs specialist equipment (or a highly trained workforce) to handle. Edge emitters are therefore more expensive than VCSELs to test. The shape and size of the optical beam can be well matched to single-mode fibre, but the assembly tolerances are much more demanding. The advantages of this design are that the optical output power can be much higher, and the wavelength of operation can be 1300 nm and 1500 nm allowing the signal to travel up to 40 km at 10Gb/s or hundreds of kilometres at lower speeds. In the Fabry-Perot laser, a simple version of the edge emitter, the laser cavity is defined by the cleft facets at either end of the chip. The output is composed of light at several discrete wavelengths, giving rise to a relatively broad spectral width. Because of chromatic dispersion (see page 8), the data stream at the output is smeared out. For longer distances, this design is not adequate and a different type of narrow-spectral-width edge emitter is needed, the distributed feedback (DFB) laser.

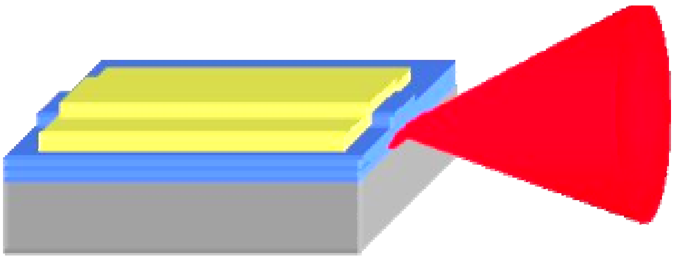

Figure 9 Schematic of an edge-emitting laser

2.2.3 DFB LASER FOR LONGER DISTANCE

In the DFB laser, a grating is built into the structure to define the operating wavelength more exactly (Figure). The light emitted is nearly monochromatic (i.e. consisting of a single wavelength), making dispersion much less of a problem. Coatings are added to each end of the structure to give the required performance. Fabrication requires more semiconductor-processing steps than Fabry-Perot laser diodes, not only to add the grating but also coatings required on the ends of the device, which can only take place after the chips have been cleft from the wafer into bars and re-processed. Furthermore, the yield associated with each of these extra steps is less than 100%. All of these complexities

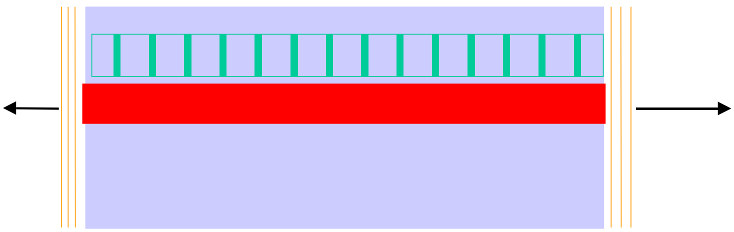

Figure 10 Schematic of a DFB edge emitter

make the DFB more costly than the Fabry-Perot laser.

As a rough guide, there is an order of magnitude difference in cost between the three types of laser described above. So if a VCSEL costs $1, Fabry-Perot edge emitters cost approximately $10 and DFB edge emitter, approximately $100. Factors such as the designed data rate of operation and the volume of manufacture also strongly influence costs, so a high-volume low-speed edge emitter can match the cost of a low-volume high-speed VCSEL.

2.2.4 Operating temperature range

The exact operating range depends on the customer's environment, the type of heat sinking designed into the equipment rack and the type of building the equipment is used in; typically the operating range of 0-70°C is defined for the equipment board. This translates to about 90°C on the optical component after heat removal (through forced or simply convective airflow across the transceiver module) and heat sinking of the mechanical design have been taken into account.

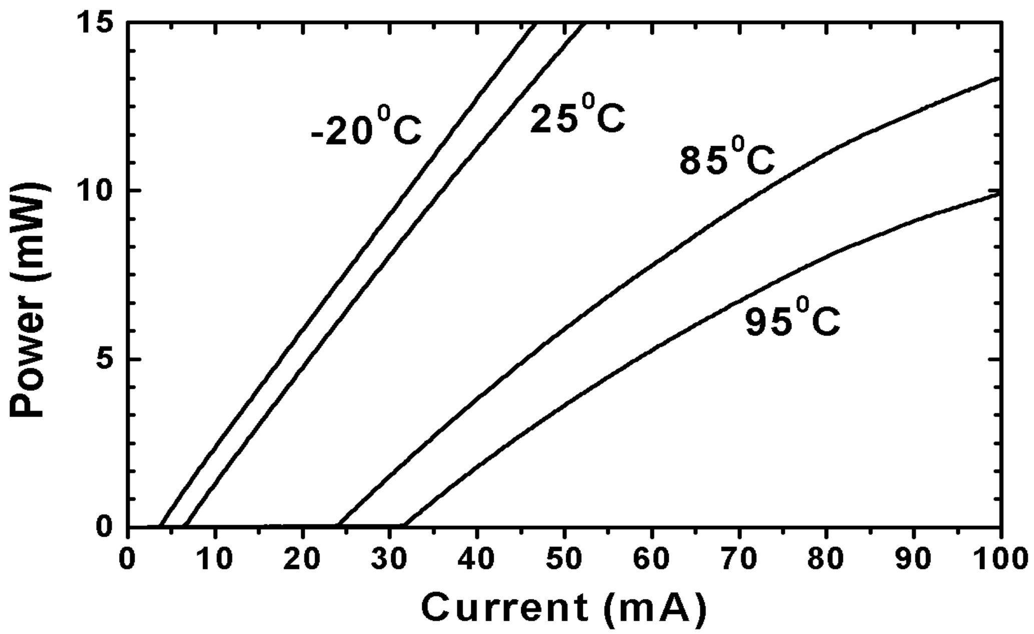

Figure 11 The variation of optical power with operating temperature. Figure used with the kind permission of Agilent Technologies

The laser is the most temperature-sensitive of all the optical components and in transmitters designed for high-value telecommunication systems, a thermoelectric cooler (TEC) is included to keep the laser temperature constant. In data communication transceivers, this option is not available. The added thermal load of the TEC and its efficiency mean that the extra heat removal would not be viable in a small footprint. The cost and space taken up by the cooler are also not viable.

As a consequence the performance demands on lasers for data communications are high and many companies are developing[6] and releasing products designed specifically for this market.

Figure 11 shows a typical light-current characteristic for an edge-emitting laser. The optical power has a threshold current before light is emitted. At -20°C this is less than 5 mA and 30 mA at 95°C. The slope efficiency also changes with temperature. At 25°C, 10 mW is achieved 30 mA above the threshold, but at 95°C only 5 mW 30 mA above threshold. The electronic laser-drive circuit compensates for the optical power changes, in both threshold and efficiency, by adjusting the drive current to minimize performance variations due to temperature changes.

2.3 Photo-detectors

Photo-detector diodes absorb the optical power from the fibre and convert it into current for electronic amplification. They are considerably simpler to design than lasers and their performance is less prone to temperature variation. The detector's active area – where the light from the fibre is absorbed – must be larger than the cross section of the incoming light beam. It is important too for the light-to-electrical efficiency at the signal wavelength to be high and, of course, for the detector to be fast enough to detect the signal. There is a choice of structures. Light can enter the chip from the top, through the edge or through the substrate. The choice is made according to the details of the product. The schematics for these different types are similar to those shown above for the laser diodes except that the light is absorbed rather than emitted.

2.4 Multiplexing

Multiplexing is the general method of combining multiple signals from their respective wires onto a smaller number of wires. By carrying the same data on fewer wires, it reduces the number of connections that must be made and makes transmission more efficient. The most common method is time multiplexing, an electrical technique, and this is used to some extent in most other communication methods too. Optical communication adds another degree of freedom; different wavelengths (or colours) can be used in addition, giving still greater efficiency.

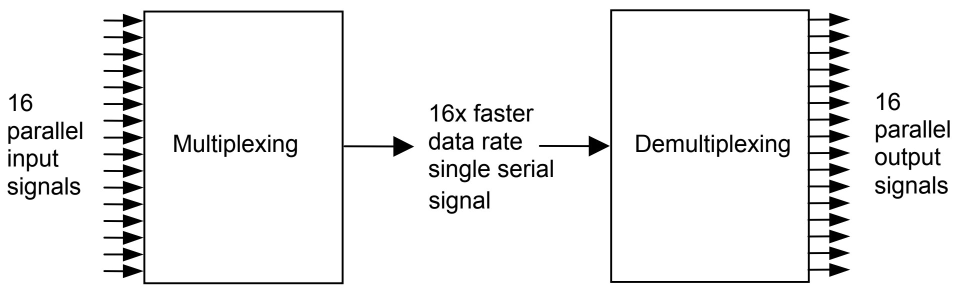

Figure 12 Schematic of time multiplexing, in this example combining 16 signals

2.4.1 TDM

Time division multiplexing, TDM, is the method of putting electronic messages together from separate parallel electronic signal to a higher-speed serial signal. At its simplest, if 16 parallel signals are combined together then the combined single output is 16 times faster. In practice, other functions are often added to a multiplexing chip (because of the ease with which multiple functions can be integrated on silicon chips. Demultiplexing from one higher-data-rate to several parallel lower-data-rate signals is the opposite process. See Figure 12.

2.4.2 SCM

Sub-carrier multiplexing or SCM is widely used in non-fibre communications; all radio systems (including, for example, those that broadcast BBC radio and terrestrial television channels) use a radio-frequency (RF) carrier and electronic multiplexing to give increased sensitivity and higher data rates. Again, this method has many variants, depending on how it is used, whether for data or text over the mobile-phone system or the wireless local area networks offering internet access at fixed-wireless hotspots such as airports. If data on RF carriers is then applied to an optical signal replacing TDM, then this is capable of transmission over greater distances on multimode fibre. Cable-television distribution over fibre uses SCM[7] techniques, but it has yet to be commercially deployed for pure data communications over fibre.

2.4.3 Parallel fibres

Parallel fibre links usually have some limited multiplexing of the signal before transmission takes place over multiple fibres. New parallel fibre ribbons of 50-µm-core multimode fibre are now available. For short links, between equipment racks within the same building or room, for example, this ribbon can be installed. It has the advantages of high data rate and low cost. Many manufacturers now offer parallel transmitter and receiver modules for these ribbon fibres.[8]

2.4.4 WDM

Wavelength-division multiplexing, WDM, is a method of combining many moderate-speed data channels using different wavelengths of light for each, without increasing the data rate of the signal.

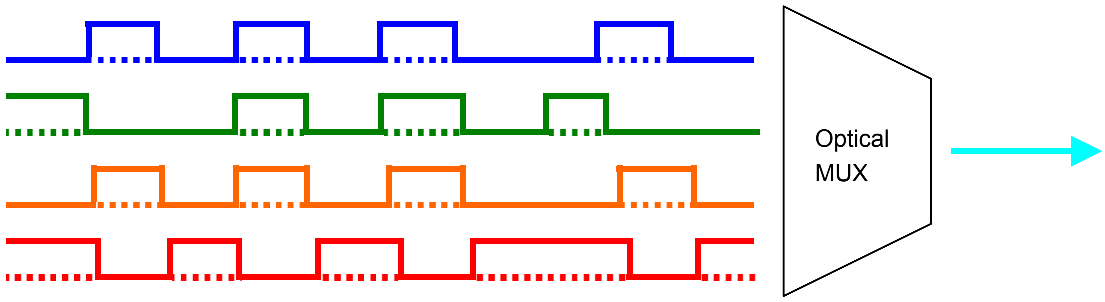

Figure 13 An optical multiplexer combines multiple signals at different wavelengths into a single multi-wavelength WDM signal

This is a significant advantage when the data rates are approaching the speed limits of electronics. Optical multiplexers (MUX) are separate optical assemblies for combining the light signals together after all the electronics multiplexing is complete (Figure 13). Complementary demultiplexer (DEMUX) devices separate the wavelengths into individual channels at the end of the fibre link, where they can then be separately detected individually. (An optical prism is an example of a demultiplexer; it can be used to separate white light into its constituent rainbow of colours.)

WDM has revolutionised the long-haul telecommunications limits over the last few years and has ensured that already-installed fibre cables can carry increasing amounts of data without further cable being installed. Commercial WDM systems carry typically 8–16 different wavelengths and research papers describe 200 channels or more, with commercial mux-demux optics now being 40-channel units. This increasing channel density is referred to as dense WDM or DWDM. The channels of a DWDM system usually have a separation of 1.6, 0.8 or even only 0.4nm, which gives a very large number of channels in the low-loss fibre window, where optical amplification is used.

A more promising type of WDM for data communications is coarse or wide WDM (CWDM or WWDM), since it will have inherently lower cost. One of the cost factors in DWDM is the need to control the temperature of each component accurately to ensure the signal does not drift into the next channel. The emission wavelength of a lasers is inherently dependant on temperature. A 115°C-temperature change corresponds to a wavelength change of about 12nm (Figure 14). A typical operating environmental temperature range for the equipment that contains the lasers is 100 °C. The corresponding change of wavelength must be accommodated. The variation would be many channel slots of a DWDM system and the laser temperature would have to be more closely controlled.

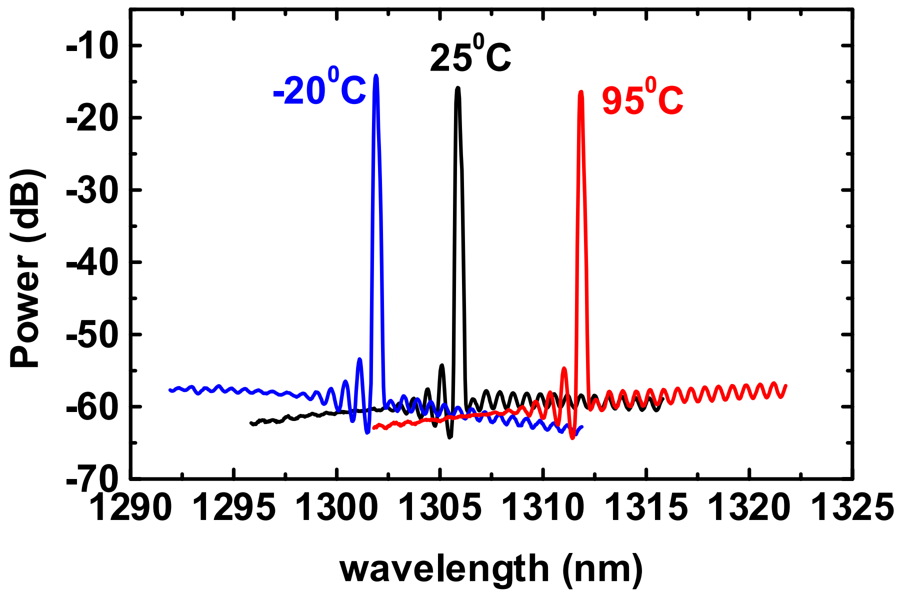

Figure 14 Inherent change of emission wavelength with laser temperature (used with the kind permission on Agilent Technologies)

The CWDM method allows the laser to drift with temperature within a 1-nm flat band and channels are placed at 20-nm intervals. This technique has been standardised in the 10-GbE standard for four channels, each of 2.5 Gb/s, around the low-dispersion 1300-nm window. It has the benefits of a lower data rate and suitability for the installed base of multimode fibre. [9] The main disadvantage of using four lasers and four respective sets of drive electronics is that they may cost more than single-channel systems operating at the combined 10-Gb/s data rate. At increasing data rates such as 40 Gb/s, when the higher-speed components cost proportionately more, or at higher speeds beyond present technology, then the CWDM method will be more popular. A simple upgrade path for higher capacity is one advantage of WDM. One fibre set is installed and later, when more signals are needed, more transmitters and receivers of a different wavelength are used without the labour-intensive stage of pulling more fibre through building ducting. CWDM is starting to be used now, with multiplexing units separated from the electronics[10] to allow each different wavelength to carry different data rates and allows future capacity upgrading.

3.0 Mechanical technology

The housings and other mechanical parts of optical communications components are important, serving as they do several independent functions at once and subject to tight cost control. The functions include optical connection, electrical connection, protection from the environment, and the enabling of optical sub-assemblies. Optical-fibre pigtails are never used in data-communication modules because the plastic-coated fibre needs special care on handling and assembly onto the PCB. All modules have push-fit optical connectors. Even screw-fit or bayonet-action optical-connector housings are too labour intensive to be used for the data-communications market.

All transceivers have push-fit optical-connector interfaces. The issues are as follows.

- Pluggability, i.e. the flexibility and self-adjusting guiding needed to ensure that connectors of slightly different dimensions are pushed together into alignment, with the necessary accuracy to ensure good optical coupling
- optical loss and reflection
- durability, such that the contacts and guide will endure for many thousands of plug-unplug cycles
- mechanical tolerance to achieve the specifications for optical loss and reflections
- latching to hold the fibre cable in place and allow easy release when needed
- cleanliness of the optical interface, which must be maintained at all times.

The electrical connections to and from the transceiver may be provided by simple circuit pins or an electrical connector. At data rates below about 1 Gb/s simple solder-once circuit pins tend to be used, but at higher data rates, ball-grid-array (BGA) connections are used instead. Electrical connectors for the data lines give the advantage of changing or adding further transceivers in service. There are many versions of both electrical and optical interfaces, from formal international standards to more de-facto versions defined by multi-source agreements (MSAs), which allow different suppliers to produce parts which work together.

The overall transceiver needs some form of case to prevent damage to internal components, ease handling and possibly provide a frame onto which the optical and electrical connectors are mounted. This case can be moulded plastic, metal (usually pressed) or metallised plastic. There is an advantage in having an electrically conducting outer case: the electromagnetic screening it affords helps contain the electromagnetic radiation emitted by high-speed electrical switching, reducing potential interference.

The optical subassembly is often a separate item which allows it to be tested before incorporation into the complete final assembly. It can provide mechanical support for the fibre with which to attach it close to the optical source or detector chip. Optical-component stability must typically be held to better than 0.5 microns for single-mode fibre products and a few microns for multimode products. This difference in tolerances is one of the main cost drivers and reasons why single-mode fibre systems are used only when data rate or distance requirements demand it. The sub-package can also provide a hermetic seal to prevent moisture ingress to the chip, which can otherwise cause performance degradation and reduce its operating life.

The full details of the mechanical design are complex. Each supplier has its own preferred techniques, which give rise to different dimensions, although all are below a maximum if fully specified. Heat-sinking design also differs from one supplier to another. Air-flow and heat-sink contacting techniques are both used. All of these issues, coupled with the exact earthing techniques used, result in differing electrical performance with respect to electromagnetic compatibility (EMC) and differing thermal performance. So, despite the intent of the multi-source agreements to ensure inter-operability of differing vendors' products, transceiver modules are not identical.

4.0 Electronics

The cost-sensitive nature of the data communications market has led over the years to a migration of the electronics technology in transceivers from a reliance on discrete transistors and amplifiers in hybrid assemblies in the past to custom ICs today. A typical transceiver can contain three individual ICs – a laser driver, a preamplifier which amplifies the photo-current to an amplitude measured in millivolts and a decision IC, which converts the analogue signal from the preamplifier to digital ones and zeros. Some manufacturers offer transceivers where these three functions are integrated onto a single chip, making cost savings in piece parts and assembly. A parallel transceiver will also contain parallel circuits to drive up to 12 parallel lasers, and these too will be arrayed to give a single IC for reduced assembly complexity and cost.

IC technology historically separated into the use of silicon – bipolar and CMOS – at data rates under 1 Gb/s and gallium arsenide (GaAs) at the higher data rates. Typical companies active in this area include Philips, Infineon, Anadigics, Vitesse, AMCC, all companies equipped with their own semiconductor manufacturing facilities. Another set of companies including Multilink and Broadcom[11] are design-and-marketing houses and use commercial foundries to fabricate their IC designs.

The speeding-up of silicon CMOS technology has brought data rates of 10 Gb/s – once the preserve of GaAs and bipolar SiGe technologies – within range. For the highest-speed ICs, the technology used is based on indium phosphide (InP), a compound-material technology similar to GaAs in design and fabrication techniques. New companies have been launched to design ICs for optical communication on InP[12,13,14] and also new foundries,[15,16] mostly for 10 and 40 Gb/s.

At higher data rates there is an increasing tendency for the multiplexing and demultiplexing chips to be included in the same primary package. This has the advantage of reducing the data rate transmitted through the electrical connections. On the downside, the electrical connector becomes larger, since it has more data connections, which may dictate the size of the package, and hence how many can be used per card. Also the module becomes application-specific, with a given data rate and format. When multiplexing is included, the modules are known as ‘transponders’.

5.0 Transceiver modules

The design of data-communication transceiver modules involves trade-offs between electrical and mechanical considerations. The fundamental performance figure is the transmission bandwidth. The higher data-transmission distances feasible with single-mode fibre (40 times greater at 10 Gb/s) are bought at the price of much tighter requirements on mechanical alignment, (50 times as accurate in the most demanding dimension). Two further examples are described in detail below in sections 5.3 (fibre coupling) and 5.4 (the electrical connector).

The choice between parallel optical links and a single serial link is instructive. If the need is to connect together two equipment racks 300m apart with approximately 30 Gb/s of bi-directional data using only commercial components, this can be achieved in any of these three ways:

- two Paroli[17] 12-wide parallel transmitters and receivers and 62.5μm ribbon fibre for each direction
- one 12-wide 30-Gb/s transmitter and receiver[18] with a 12-wide 50-μm ribbon fibre for each direction
- three 10-Gb/s transceivers with three duplex fibre links (either new-design multimode fibre or single-mode)

The most compact and probably lowest-cost choice of these is the single 30-Gb/s transmitter and receiver. If the fibre distance were longer or the equipment designer wanted a common interface for any distance, only the 10-GbE (10-Gb/s Ethernet) single-mode option would suit.

On a typical communications-switch circuit board the optical transmitters and receivers occupy a very significant part of the board area (the footprint), require external panel space for accessibility to the optical connectors and can cost as much as 80% of the total board cost. This means that transceiver-module manufacturers are under continuous pressure to reduce PCB footprint, the overall cost of the modules and to provide compact optical connectors.

The concept of hot-pluggable modules has been a solution to these demands. Transceivers are designed with both optical and electrical connectors so that no special handling of fibre is required and the module does not have to be soldered into place. The transceiver can be added to the PCB whilst it is powered and in use, hence the ‘hot’ prefix. The transceiver connector mates with another connector provided within a cage on the PCB, allowing the expensive optical transceiver to be added only as required. When a transceiver is hot pluggable, a short-reach transceiver designed for multimode fibre can be replaced with a single-mode fibre and transceiver should a long transmission distance be subsequently required. This has led to families of transceivers where a manufacturer designs many data rates and types of transceiver, which are the same size and have identical pin connections in the same electrical and optical connector. The advantages for the equipment manufacturer are just-in-time population of the PCB, custom configuration for each user and future expansion of service in the field, all giving cost savings.

5.1 Assembly

Optical assembly is the most critical part of the process. This can be the direct alignment of an optical fibre with the light source or coupling via an optical lens. The placement accuracy needed is expressed in microns (μm), (1000 μm make 1 mm). For a surface-emitting laser and a multimode fibre the accuracy required is approximately 10 μm in the x and y dimensions and approximately 40 μm in z dimension (the distance between the emitter and the laser). The exact distances depend on many factors such as the size of the emitting area and the angle of the light beam. The most efficient optical coupling will take place when the shape of the optical beam is the same as that of the fibre core, which is circular. If an optical lens is used, this must be placed with similar accuracy but the placement of the fibre is more relaxed, especially in the z dimension.

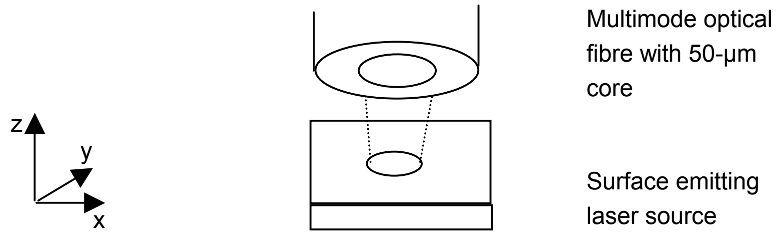

Figure 15 Optical fibre being manipulated above the laser to give the maximum optical coupling

Edge-emitting lasers do not have this inherent circular emitting area but, more often, a rectangular one with dimensions typically of about 1μm x 2 μm. This must be coupled to the 7-μm round core single-mode fibre. Despite the mismatch, coupling efficiencies of greater than 80% (defined as the ratio of the light power emitted from the laser facet to the light power injected into the fibre core) can be achieved with optimised laser and lens combinations. The placement accuracy for the optical component in order to achieve high coupling efficiency can be much less than 0.5 μm in one key dimension. As well as sub-micron accuracy demanded for component placement, it must remain in place throughout its expected operating life and survive the necessary qualification stages. A high coupling-efficiency optical system could contain three lens elements, each requiring to be fixed with high accuracy; a movement in any single component could cause unacceptable changes in the optical power of the final optical subassembly.

At high data rates, reflections from interfaces in the fibre can cause unacceptable changes in the performance of the transmitting laser. This necessitates the coupling of an optical isolator (to block reflected light) to the laser. If the optical coupling efficiency were 80%, then 64% of reflected light is coupled back into the laser; but at only 20% coupling, a mere 4% of reflected light is coupled, giving effective isolation. Higher optical coupling efficiencies require greater isolation and lenses to focus the light through the isolator. Higher data rates and higher-power transmitters thus call for significantly more complex optical coupling schemes.

Detectors have more optical coupling tolerances in the region of +/-20 μm allowing simple lenses or direct fibre coupling to be used. However, as data rates increase, the size of the detector must be reduced in order to reduce the capacitance. The tolerances are then similarly reduced in order that all the light from the fibre is absorbed by the active area of the detector. Each manufacturer uses techniques established in-house for laser coupling.

When coupling parallel modules to ribbon fibre, there are advantages in using arrays of optical components, since 12 individual optical alignment stages are replaced by one. This has been demonstrated by research at HP labs[19] indicating the benefits of arrays. The design of packaging for parallel components to achieve considerably lower costs than 12 individual transceivers is also crucial. Infineon[20] have reported techniques of packaging parallel components.

For less precise attachment of optical components, epoxy adhesives can be used. An optical manipulator is used to place the fibre whilst monitoring the coupled power and the epoxy is cured in situ with UV light. Solder attachment of metallised optical fibre has also been successfully used for permanent fixing. It is however laser welding[21] that has emerged as the most popular technology for attaching single-mode components. An intense laser beam focused onto the two metal parts to be fixed is pulsed to melt the metals; when cool, the joint is welded together.

5.1.1 Hermeticity

The optical subassembly (OSA) is most commonly packaged hermetically to prevent moisture ingress to the semiconductor material and dust accumulation in the optical pathway. This final lid seal of the OSA is frequently welded but can also be soldered. The fibre interface is usually the most challenging to seal hermetically. Many approaches have been used. Fibres are frequently metallised where the glass-to-metal seal is chemically based (and hence both strong and leak-free), and this metallised fibre is then inserted into a metal tube so it can be soldered or welded into the package wall. Glass seals are also used to seal fibres into metal tubes that can be soldered or welded into the sidewall. Where coupling efficiency is not critical, lensed packages (in which the lens forms a glass-to-metal seal) can be used to couple to an external non-hermetic fibre assembly. If hermeticity is critical to the reliability of the component, leak testing must be carried out to verify the seals.

Non-hermetic packaging has also been used in data communications.[22] In this approach, special coatings are critical to keeping moisture from causing degradation to the semiconductors and the optical pathway.

The de facto standard for reliability and performance assessment of optical components is that of Telcordia (previously known as Bellcore), a US telecommunications service provider. Specifications exist for many, but not all, components. The type of conformance and environmental and life tests which must be passed to prove suitability for use are defined. Since some components are not specified and the testing requires large numbers of components to be destructively tested over a long period of testing, the phrases ‘meets the intentions of the Telecordia specification’ and ‘Telcordia testing is underway’ are frequently seen. For large-volume products, the test and qualification overhead is affordable but for smaller-volume designs, it can result in higher prices and be economically unrealistic.

Both hermetic and non-hermetic packaging methods have passed the Telcordia qualification and reliability requirements for optical components, which require endurance testing at high temperature and again with temperature and humidity simultaneously high.

5.1.2 Electronics assembly

The electronics assembly mainly uses standard processes. More than five years ago, this was bare chip and wire on thick-film hybrids in gold-plated metal hermetic packages. For lower-cost assembly, surface-mount pre-packaged ICs are used, or unpackaged bare chips are placed and wire-bonded onto standard PCB material. To give protection from the environment, the bare chips are covered with industry-standard encapsulants such as Hysol, [23] which is designed for chip-on-board production. The supporting circuit components of resistors and capacitors are solder-attached with standard surface-mount techniques.

The electronics assembly methods are determined by the standard processes used within each company and are mainly decided on the cost and performance criteria for the particular product under design.

5.2 Assembly equipment

In the optoelectronics industry, many components are custom-designed for individual equipment suppliers. What's more, because of the pace of change, production runs have often been fairly short. Partial automation of the assembly process is nevertheless common, especially in the form of software optimisation of predominantly manual processes.

Compared with the semiconductor industry, automation of assembly has not developed as intensely, since automatic equipment is not as cost effective in the optoelectronics industry. But this is expected to change so that unit costs can be cut in future. The Optoelectronics Industry Development Association (OIDA)[24] holds frequent technology workshops to 'develop and refine a North American consensus for a five-to-ten-year roadmap' on the pattern of the influential Sematech alliance in the silicon industry.

The data-communication sector of the optoelectronics industry is further along the path of cost reductions when compared with long-haul telecommunications components. The standards process together with multi-source agreements (MSAs) allows both common mechanical and electrical specifications to be defined, which allow in turn high volumes to be generated in a common form factor. Automated assembly[25] reduces the labour content in manufacturing cost and is generally associated moreover with greater uniformity in product performance. The degree to which automation is cost effective depends on the unit cost of labour used, which is different in the West (where high-end optoelectronics for telecommunications are commonly manufactured) and the Far East (where data-communications modules are often manufactured).

5.3 Fibre coupling chip-mechanical trade-offs

Fibre coupling tolerances for single-mode fibre can be eased by introducing laser sources with optical modes closer in size to those of the optical fibre. This design is referred to as a spot-size converter and yields a significant relaxation of the sub-micron assembly tolerances normally associated with single-mode coupling. This can be achieved during the laser manufacturing process (in wafer form) by incorporating modifications to the optical waveguide structure. Typically this introduces significantly more process steps in wafer fabrication and is not yet widely incorporated in the manufacture of low-cost components owing to the extra cost and complexity needed at the wafer-fabrication stage. The size of the laser chip is also increased because of the details of the optical design, which again increases the cost because fewer chips are obtained from each wafer. When this technology becomes more mature it is expected to reduce the cost of coupling sufficient optical power significantly. In this way, the increased cost and complexity of the optoelectronic chip will be offset by a reduction in the cost of assembly.

5.4 Module electrical connections

The MSA for 10-Gb/s small form-factor pluggable (XFP) transceiver modules defines a module with serial data transfer and no electronic mux-demux IC. This has the advantage that the module is not application-specific and hence should be more widely used. The module is hot-pluggable and the transceiver module has a smaller footprint on the PCB. The disadvantage is that 10Gb/s data must be transferred through the connector. Traditional connectors for this data rate are coaxial in style and fairly expensive because they are fabricated by metal machining. A small-form-factor push-fit connector has become widely used in lower-speed (to 2Gb/s) hot-pluggable modules and consists of a host-board housing with sprung fingers push-fitted onto controlled-impedance tracks on the rear of the transceiver module. The connector is enclosed in a metal cage to ensure radiative emission is kept low. The XFP[26] connector has been selected to transmit 10Gb/s and uses differential signals (data and its complement, data-bar) with adjacent ground contacts. The mechanical constraints on the connector are durability, the mechanical force needed and spring tension. Since the connector is designed to be hot-pluggable, the connections must be made in a particular order: first ground then power supplies followed by the signal. The finger contacts must mate with the controlled impedance tracks on the transceiver PCB and give a repeatable bandwidth transmission with low electrical reflection, all of which imply a requirement for tight control of the impedance of the connector. The external cage screen to reduce electromagnetic interference has internal electrically-conducting gaskets around the transceiver and the rear connector. These gaskets must be mechanically compliant yet must also establish electrical contact. The cage also conducts heat away from the transceiver, improving electro-optic performance by reducing the operating temperature of the active components.

At the higher data rates, two signal lines are used, one using data and one carrying the exact inverse of the data. This prevents the signal from being radiated and causing interference to adjacent equipment or components. The track lengths of these two signal lines have to be virtually identical too and, as the data rate increases, the tolerances on the lengths and the amplitude matching (how exact a match the data is to its inverse) are reduced. At the higher data rates, the material of which the PCB is made also becomes critical. No longer suitable, the standard low-cost material has to be replaced by more expensive, specially designed materials.

6.0 Market

Most equipment customers and suppliers of optical data-communication transceivers have headquarters and design centres in the USA and in the standards-making processes US companies make up the highest proportion of participants. They share the market, however, with a wide array of international competitors.

The data-communications sector has been affected by the telecommunications downturn during 2001 and 2002 but some companies' financial results (where enough detail in the financial statements are available) show they have since returned to modest revenue growth, if only limited profit. All participants describe the current market conditions as difficult. And although some report signs of recovery, uncertainty remains strong in the market for telecoms optical components.[27] Forecasts from Freedonia on fibre optics[28] are for similar revenue growth rates for telecommunications and data networks by 2006. Historically the growth rates have been 15-20% with growth of 50-100% common in the telecommunications boom of 2000, and catering for the continued growth at the higher rates has led to the present over-capacity in many companies.

The product life cycle is frequently less than five years for a transceiver, and design upgrades and new product designs are underway in all the main companies to create the next generation of components. Opportunities for fresh suppliers and competitors are generated by the re-design cycle, especially if there is a potential for both cost reduction and performance improvements. Each manufacturer must demonstrate reliability and conformance of the equipment and the components and therefore there is a qualification stage for each product. Having demonstrated compliance companies are reluctant to change supplier since this would require a re-qualification. This leads to hurdles of entry to fresh products and ensuring design-in of a component is critical to commercial success, hence timing of the component product launch is crucial since it must be available for the customer to verify performance and complete equipment qualification using that part.

The industry is restructuring so that risks are shared. Assembly is often outsourced, leading to opportunities for contract assembly companies. The high overhead of supporting a dedicated large manufacturing capacity to cater for the peaks of demand is difficult to sustain in the troughs; therefore contract assembly, where the demand is smoothed over many industrial sectors, is mutually advantageous. Volume manufacture tends to be partially automated and take place in economies where labour costs are low, typically in the Far East or other developing economies. Equipment manufacturers also have production facilities or manufacturing subcontractors in lower-cost countries.

Ethernet is the clearly dominant data-communication protocol, and for backbones in buildings, where optoelectronic solutions are essential, GbE dominates. 10 GbE is the most recent opportunity as data rates increase still further. The use of common standards, which ensure components from different suppliers will work together is a dominant part of the marketplace, and also ensures that the cost of the products is a crucial selling point once design-in has been established. In other words, this is virtually a commodity marketplace. Competitors vie on price and availability and, beyond that, only in the race to satisfy new standards or introduce new products.

7.0 Standards

It is not possible to describe here every optical specification or standard. Instead, a selection of the more important ones is introduced. The International Telecommunications Union (ITU) is responsible for many optical communications specifications which are also applied in data communications. In general the ITU specifications are most appropriate for the longer-distance high-performance links, where manufacturers' equipment does not need to work together and where price competition is not so dominant. On many occasions a standard which started in the data-communications arena has also been standardised in the ITU. The specifications highlighted are the dominant ones in fibre data transmission or are, in the author's opinion, those most relevant to future opportunities.

Standards bodies have many forms. For instance the ITU has national representatives from each member state, who take instructions from their respective national regulatory bodies. In principle, each country advances a view representing the interests of all interested parties in that country, from industrialists to licensing authorities. It is a legalistic and time-consuming process. Making any changes is expensive. In comparison, the OIF is industrially biased and faster moving. Its objective is to increase the market acceptance of particular products. SMEs can benefit greatly from contributing to standards development in bodies such as the IEEE or the OIF (with some transatlantic travel costs), since it ensures that suppliers know what their customers want and can influence the outcome of the process. Attendance at a standards-body meeting can raise the profile of an SME and generate more sales leads than attendance at a trade show. This is because technical competence is demonstrated, discussions with customers can be arranged as side meetings, the activity of competitors can be monitored and activity is more focused than at most trade shows.

7.1 GbE

This Ethernet standard with a data rate of 1 Gb/s was first proposed in late 1995 and was completed and approved by the Institute of Electrical and Electronic Engineers (IEEE) in the USA in July 1998. Ethernet standards are always specified to work worst case on the links and the general design principle is that they should always work. The detailed technical understanding which goes into the design of the specification ensures that different vendors' components are compatible and has contributed to the commercial success of the specification. A detailed book explaining the technology and giving a good insight into the process of standardisation is in the further reading list (p. 36).[29] Bristol University had a significant role in solving some of the technical issues raised. Other texts give more usage and application areas. The specification as issued by the IEEE is available[30] on the web. The volume of Gb/s optical transceivers that have been sold is considerable, some two million units per year. The optical variants include a single-mode version to cover the longer distance as well as multimode versions at wavelengths of 850 nm and 1300 nm for the shorter-reach optical links.

Ethernet has become a self perpetuating success. The volumes of components which are designed for Ethernet are much higher than for other standards or for proprietary designs and costs are concomitantly lower.

7.2 10 GbE

Inspired by the success of the GbE standard, 10 GbE is 10 times as fast. The standards process was started in March 1999 and the formal ballot approval was in July 2002. The issued standard is now available from the IEEE. The physical layer, which defines the actual hardware and items used, defines five different distances, according to the fibre type used. Installed multimode fibre has the lowest bandwidth and hence the shortest transmission distance, 100 m. Newer designs of multimode fibre have higher bandwidth and hence allow link lengths of up to 300 m. Beyond 300 m only single-mode fibre is defined, with three specified lengths of 2 km, 10 km and 40 km. The latter is an innovation in Ethernet terms since this distance is considerably longer than in previous specifications.

The Ethernet standards process does not determine the physical dimensions or footprint of the modules designed to be compliant, simply because at this level of detail it is difficult to reach agreement within the time frame. The physical dimensions are determined by multi-source agreements between competing vendors. Because they design for different applications and different customers, variants emerge of package style and interface details. Information about these MSAs and the partners supporting the platforms can be found on their web sites – for Xenpak,[31] XFP,[32] Xpak[33] and X2.[34] All of these modules are physically smaller than the existing 10-Gb/s modules, which were designed more for telecommunications use under what is generally called the 300-pin MSA (a name alluding to the type of electrical connector specified).

Opportunities for new entrants and significant revenue growth for existing players are higher with this recently agreed standard. There is a good deal of activity between competitors on MSAs and module platforms.

7.3 CWDM

Coarse wavelength-division multiplexing (CWDM) was first proposed for data communications. It has the advantage that the data can be multiplexed together at a lower data rate using different wavelengths of light. One option in the 10-GbE standard is for four channels of 2.5-Gb/s data with four wavelengths all in the 1300-nm window.[35] In the 10-GbE implementation, the module uses four lasers, four photodetectors and electronics plus the optical mux-demux, but unlike the 10-Gb/s serial version, it does not need the 4:1 electrical mux–demux. One significant performance advantage of the CWDM approach is that it can be used for both multi-mode and single-mode fibre and is the only standardised technical approach that can reach transmission distances of 300 m on the older, already-installed multi-mode fibre.

CWDM is significantly less expensive than dense WDM (see page 14), because the temperature does not have to be controlled. The approach has also been standardised as ITU G694.2 with the same 20-nm channels across the whole optical band. A recent paper[36] reported use of a fibre with a suppressed water-absorption peak (cf. Figure 3) and by using the central wavelengths too, a data rate of 160 Gb/s was achieved over 12 km using 16-channel CWDM. This technique has this impressive capacity only in combination with the new fibre, but this would be a cost advantage in freshly-installed fibre cables.

7.4 OIF

The Optical Internetworking Forum[37] has as its objective 'to foster the development and deployment of interoperable products and services for data switching and routing using optical networking technologies'. It is a predominantly industrial forum that aids the development of the standards. Previous similar forums have included the GbE forum, which supported the development of the GbE standard. The forum allows the development of technical consensus between potentially competing technical approaches from vendors. The forum uses existing standards where they are available and can provide additional detail to supplement these existing standards where appropriate as well as assisting in the development of new standards.

The aims of each process always include targets for cost of the components in relative terms, for instance where a 40-Gb/s standard is to be developed, the choice of technology is determined by those most likely to be less than two-and-a-half times the cost of a 10-Gb/s link. Clearly this economic target accentuates the further commercial pressure to reach agreement on the standard.

7.5 FSAN

FSAN – Full Service Access Network[38]– is a standard that was started by telecommunications service providers to offer lower-cost fibre connections to business and domestic customers. In each country the existing type of equipment and the strategy for deploying access networks is different. In order to achieve cost reductions through higher volumes (perhaps tens of millions of access lines), many equipment suppliers, component suppliers and service providers must offer competing interchangeable products. This required international co-operation in FSAN and when it was mature the standard was ratified formally within ITU standard G983.

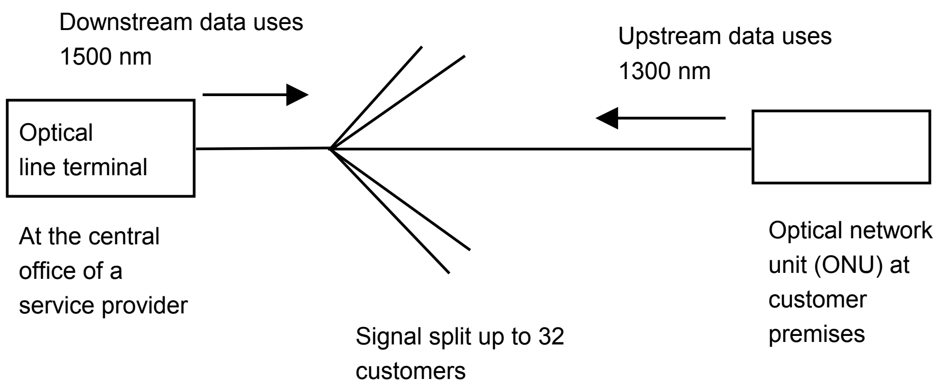

Figure 16 Schematic of a fibre system using the FSAN standard

The basis of FSAN is a passive optical network (PON), which is simply a power-splitting and service-sharing method, as shown in Figure 16. The signal to the customer is transmitted using a more expensive style of 1500-nm laser, but the cost of the source is shared between up to 32 subscribers. Each home or business premise has a lower-cost 1300-nm laser (and significant electronics) installed when they subscribe to the service. The customer's upstream signal does not interfere with the downstream one even though it is on the same fibre because it is on a different wavelength. Several customers may try to communicate upstream at the same time and these signals are prevented from interfering with each other by a simple time-sharing method of taking turns. The standard has been developed to take data and telephone services and distributed video too. PONs have also been installed by cable-TV operators, using the same PON shared approach but each has different electronic methods.

The economic return from a PON installation depends strongly on the proportion of subscribers who take up the service. If only one subscriber out of the 32 potential customers takes the service, the installation and equipment costs are shared between fewer customers. However, if the price of the service can be reduced to attract more customers once the fibre is installed, the potential could be great since most of the costs are by then sunk, both financially and literally into the underground conduits. For higher subscriber take-up rates, the PON has commercial advantages but for lower take-up a dedicated fibre is no more expensive.

7.6 EFM

EFM is Ethernet for the first mile, that is the first mile from the customer premises, previously referred to as FTTH (fibre to the home). This latest Ethernet initiative from the IEEE is building on the commercial success of previous generations of the standard, which have been the lowest-cost data-communications choices. Supporters of FSAN acknowledge that this next generation has the potential to reduce costs further. The EFM standard has several proposed physical implementations mostly catering for distances of up to 10 km, with one 20-km option. This allows the fibre to go indirectly from the central distribution point to the customer, but it is clearly anticipating distances of more than a mile. The fibre use is only single-mode, which is a direct consequence of the distances needed to be covered. The data rates selected are 1 Gb/s and 100 Mb/s. Developments can be monitored at the EFM web site,[39] where the open technical issues and presentations given by each attendee to the standards meeting can be viewed.

The differences between the proposed options are in the type of fibre configuration. One option is a single fibre. This needs to carry data in both directions and the link must be designed to prevent mutual interference. The proposed method is to use two different wavelengths – one upstream in the 1300-nm window and one downstream in the 1500-nm window. Wavelength is also used to split the upstream and downstream directions in the passive optical network configuration (PON), which has a potential split ratio of 1:16, thus sharing the cost of the downstream laser and equipment between 16 subscribers.

The dual-fibre systems have a fibre in each direction for upstream and downstream traffic, which although it clearly uses twice as much fibre (which costs more), allows simpler optics (which cost less) than single-fibre systems. The cost of installing the fibre into the ground is common for all the fibre approaches.

8.0 Future

There are three general but dominant technology trends in transceiver modules, namely size, data rate and cost. The size of modules is decreasing as more integration of circuit technology is used and optical components are built in arrays. Data rates are increasing with time as demands for more throughput increase and, in parallel with these two trends, the price continues to fall; each generation provides a lower cost per gigabit transmitted even when newly introduced.

In general terms, when connecting office-based PCs together there is no fundamental limit to the data rate. The most common connection speeds at present are 1 Mb/s or 10 Mb/s using Ethernet over copper cable. More complex electrical modulators will allow copper wire to be used to 1 Gb/s for short distances. Eventually fibre and components will become cheap enough and data rates high enough that fibre-to-the-desk will be common using multimode fibre. Fibre installed throughout the office tends to be multimode fibre because the larger fibre core means connector tolerances are less critical. Recent work at Essex and Cambridge universities in partnership with Fujitsu has shown that multimode fibre can be used to transmit up to 1000 Gb/s over 3 km.[40] This demonstrates that if there is a large enough demand for bandwidth, the limits are not fundamental physics but product engineering and cost. This 1000 Gb/s is the present level of traffic on the US internet backbone,[41] so it will probably be some considerable time before this is needed to each PC.

The OIDA's five-year road map shows more design-for-manufacturing techniques being used, more optical integration, more standardised packaging, more automated assembly, and considerably more contract manufacturing. Software tools to integrate the whole of the design and manufacturing process, which integrate electrical, optical, thermal and mechanical issues, are also needed. At present, software exists for each of the design aspects, which must then be integrated by experienced engineers to complete the whole design.

8.1 Fibre to the home

Fibre to the home has been the subject of many field trials around the world over the last decade or more but has not yet happened to any significant extent other than in Japan, since the cost of the connection has not been competitive with the existing links operating at much lower data rates. Some of these issues are regulatory, since it is not economically viable for each service provider, utility, telephone or cable-TV company to independently install fibre to every home (nor would planning authorities allow the repeated disruption of the pavement and road access). No satisfactory sharing of a fibre resource has yet been proposed in the regulatory environment. The other limitation is the provision of emergency telephone access in the event of mains power failure. (Ordinary domestic telephones obtain power supplied over the electrical wiring from the exchange to the telephone socket in the house and so can be used even when mains power to the premises has been cut.) Battery back-up for these conditions is viable but expensive if the equipment is power hungry, and the batteries required are large and expensive and have to be regularly recharged.

The technology for providing fibre to the home has suffered from a chicken and egg situation. The parts are not low enough in cost, yet there are huge potential cost savings if the manufacturing volumes were high enough. The breakthrough in this direction is seen as standardisation. The FSAN standard (see page 27) was designed to allow common optical modules to be designed by several manufacturers for PON networks and the latest initiative is EFM (see page 28).

8.1.1 MARKET

The FTTH market is most mature in Japan, where installations are mostly PON shared-fibre architectures using the FSAN standard or Ethernet optical point-to-point links. The following information was presented by NTT in a seminar organised by the EU's R&D support agency for the Information Society Technologies (IST) action in the Sixth Framework Programme.[42] Figure 17 shows the deployment of broadband services in Japan, which has the highest broadband take-up. Presumably other countries will follow this trend if the same conditions apply. Though penetration has only so far reached about one tenth that of ADSL, it is growing faster than both alternatives – ADSL and cable. One factor which has increased the take-up rate in Japan is the steady decrease in price. The data rate provided is 64 kb/s for cable, between 0.5 and 12 Mb/s for ADSL and 10 or 100 Mb/s for FTTH. When customers have upgraded their access method, the data they download has tripled in volume.

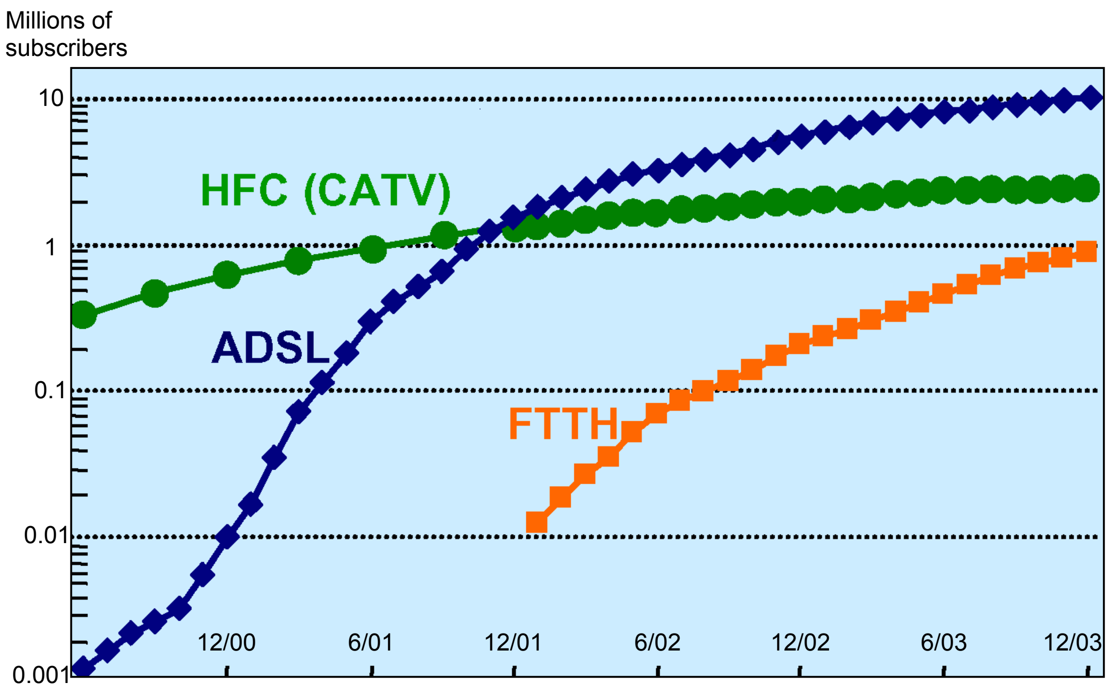

Figure 17 Growth of fibre to the home (FTTH) compared with ADSL and cable in Japan (with kind permission of NTT)

8.2 Microprocessor interconnections, chip to chip

The most common method of interconnecting chips within an equipment rack is by electrical metal tracks and electrical connectors on the PCB and by way of the electrical backplane between boards. As processors' clock speeds and data rates increase more connections at higher speed must be made and this can lead to an electrical bottleneck. Optical connections between boards on the same or adjacent shelves in the equipment rack offer a state-of-the-art solution to this connection bottleneck. To keep the optical fibres well controlled and easy to assemble, the latest commercial technique is optical-fibre flexfoils (Figure 19), in which the optical-fibre harness is pre-shaped and encapsulated in Kapton. This allows three-dimensional assembly of the optical pathway with a single optical component. This particular example allows each point to communicate with all others in a compact assembly.

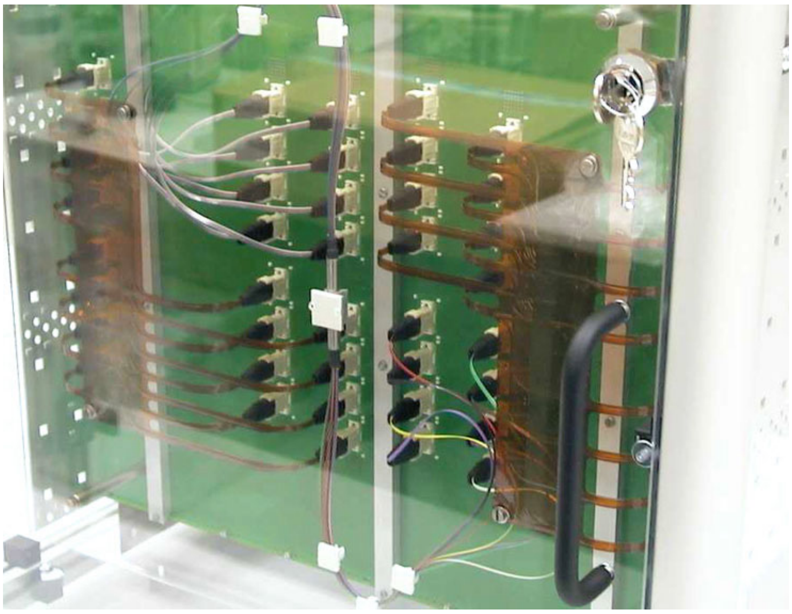

Figure 18 Optical backplane with RCI Flexfoil and connectors from FCI (with kind permission of RCi[43] and FCI)

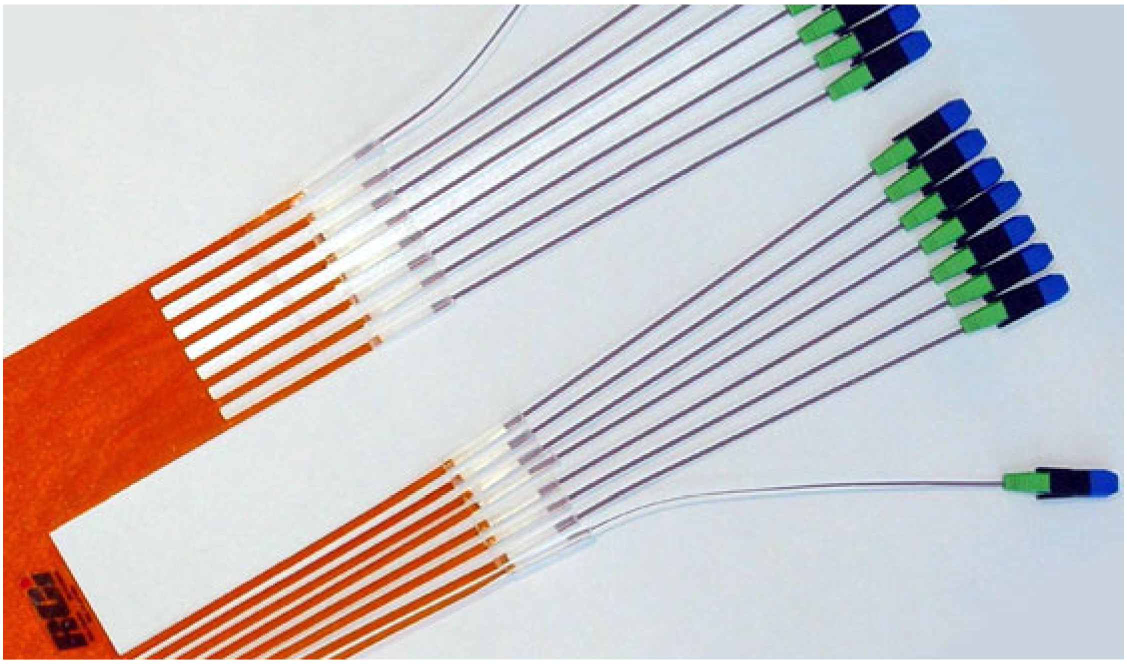

Figure 19 RCi flexfoil with ribbons and fusion splicing; connectors are MPO connectors from FCI) (with kind permission of RCi[43] and FCI)

This would be used in conjunction with optical transceiver modules on each PCB and offers the advantage that the fibre loom does not have to be routed via the backplane and connectors (Figure 18). The flexfoil shares the advantages of a flexible PCB: it allows tracking to be routed from board to board through three dimensions and at any arbitrary angle. There is no need for separate fibre links and connectors, wiring harness or cable ties, so the connection process is simplified and partially automated and cost is reduced. In the future, optical layers in the PCB may replace this flexfoil, which could be embedded fibre or glass layers as part of the PCB structure. Alternatively fibres from one microprocessor may be routed directly to another (see section 8.2.2 Technology, p. 32).

8.2.1 MARKET

There is a widely accepted view that both the backplane, for transmitting data between cards in an electronic equipment rack, and inter-chip connections on a PCB are likely to become the sites of future bottlenecks. There is less general acceptance on what the best solution is. But the highly respected and often conservative International Technology Roadmap for Semiconductors (ITRS) has stated 'Optical interconnects are considered a possible option for replacing the conductor/dielectric system for global interconnects'.[44] The ITRS road map process for IC interconnect is the pre-competitive joint view across five regions of the world and envisages the first commercial chips including optics in 2011–2012.[45] The necessary enabling research identified in this analysis is that of the optical technology and the modelling tools to integrate complete optical links into the IC and PCB design.

8.2.2 TECHNOLOGY

Research into chip-to-chip optical interconnect is very active and it is useful to describe two current EU-funded projects, IO[46] and HOLM.[47] The IO project uses two-dimensional arrays of VCSELs and photodetector arrays to connect CMOS ICs for use in IP routers. Arrays of plastic optical fibres are used between boards, and glass sheets embedded in the PCB itself, for on-board connections. Electronic design automation and integration with existing electronic-design-automation (EDA) tools is one of the key tasks for this project.

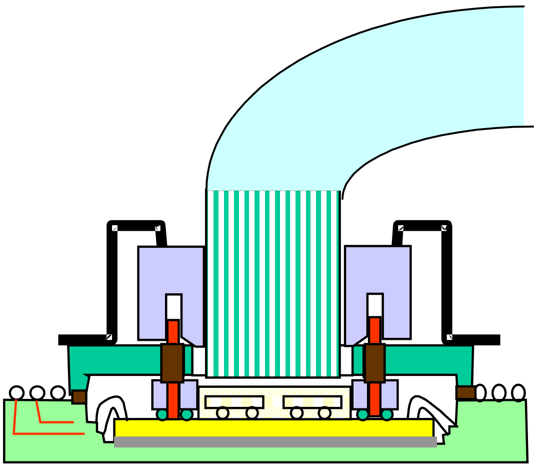

Figure 20 shows a large, wire-bonded, core CMOS processor on which the optoelectronic interface IC arrays are mounted using solder-bump technology. The optical-fibre connector bundle, with many fibres, is aligned with the optical components using pins which connect from the CMOS IC to the optical assembly. The whole component is in a surface-mount package, which interfaces with the optical connector and is soldered to the PCB. More information on the achievements of the project and other activities in the field can be found on the website.[48]

Figure 20 The concept of fibre attaching to each CMOS processor (with kind permission of the IMEC and the IO project).

The HOLM project, led by Heriot-Watt University, is developing a high-speed optoelectronic memory system using optoelectronics to connect multiprocessor CMOC ICs to provide a digital signal-processing function. The optical connections reduce the time delay between each chip, reducing memory access time. The technology approaches being investigated are free-space optics and waveguides integrated into the PCB. The university has also worked on optical interconnect for avionics applications.[49]

9.0 SMEs – where are they in the supply chain?

The industry sector is dominated by a few large players and users – among them, Agilent, Infineon (for modules) and Cisco (for routers). However there are many smaller players, of which a few exemplars are described below. It is not unusual for SMEs in this sector to possess key enabling intellectual property (IP), whether originating in the SME itself or from collaboration with a university or research organisation. SMEs frequently show more dynamic exploitation of the technology, since they do not have the structures and organisational control processes that can slow down a larger company. However, in the volume business, economies of scale for manufacturing are commonplace, and a smaller company may find itself at a competitive disadvantage. The one area with few SMEs is that of glass-fibre manufacture, which is capital-intensive and has large economies of scale and considerable invested IP. It is dominated by large suppliers such as Corning.

Module manufacturers

There are a number of module manufacturers, funded by venture capital, who are in direct competition with the existing players. These start-ups tend to be formed either from teams who are already very experienced in the industry[50] or from research institutes[51] with key enabling technology. There are opportunities with new product ranges that offer higher capacity through more parallel channels,[52] or more wavelengths.[53]

Sources

A 1300-nm VCSEL source is seen as a key enabler in data communications. The surface-emitting laser allows simple fibre coupling and on-wafer test, while emission at 1300 nm and the narrow-spectrum linewidth allow operation in the highest-bandwidth window of the fibre. This gives a higher data rate over the same distance or longer-distance operation at the same data rate. The existing technology is edge emitters and there may be savings in total module cost due to fibre alignment. There are many small start-up companies, especially in the USA, working to a goal of VCSELs and transceivers using these surface emitters. In Germany there is a spinout from the Technical University in Munich, VertiLas, which is developing 1300-nm and 1500-nm VCSELs and a laser which has the capability to tune 2 nm in the 1600–2000-nm region.[54]

Electronics

Design of integrated circuits specifically for optical communication is a key knowledge area that has generated a number of new companies. Broadcom[55] in the USA is now well established and no longer an SME, and, in the UK, Microcosm has products up to 1 Gb/s and a further offshoot company, Phyworks,[56] is offering 10-Gb/s chips. These companies are design houses and outsource manufacturing to commercial CMOS silicon foundries.

Optical assemblies

One example of an optical subassembly is an array of optical lenses for the production of parallel optical transceivers,[57] which, when built into an optical connector (Figure 21), allows all the critical optical connections to be completed in one step. The detail in Figure 21 shows the array of lenses that will couple to an array of laser emitters. DOC, a US company with expertise in the design of integrated micro-optical subassemblies, has the technology to fabricate the lens array and offer this as an integrated solution with a parallel optical connector.

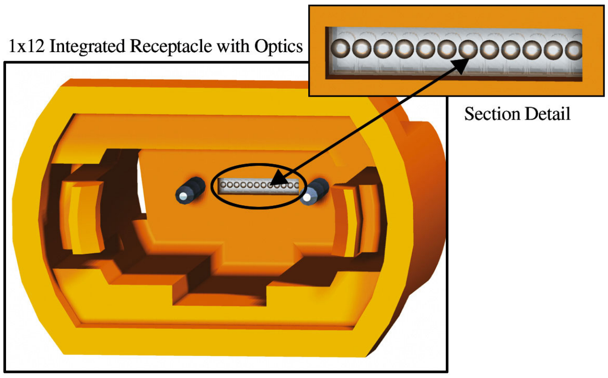

Figure 21 Lens array in a parallel optical connector (with kind permission of DOC)

Subsystems

Across Europe and the USA, several companies – including British companies such as Optical Antenna Solutions,[58] a spin-off from the University of Warwick – offer systems and components for optical communication without fibre, using focused free-space beams between vehicle and roadside as well as between buildings, where they remove the need for street-level ducting. This can be advantageous in cities where the cable ducts are very often at capacity and a free-space solution is quicker to install than digging up pavements. Direct-beam optics are also more secure than radio. This technology may prove valuable for fibre to the home too, because it obviates the need for the labour-intensive and disruptive task of digging holes for fibre ducts.

Further reading

Fundamentals

1. Personick, Stewart D. 2000 'Evolving toward the next generation Internet, Challenges in the Path Forward', *IEEE Communications Magazine*, Vol. 40, No. 7, July 2002 (Special Issue on The Internet Past Present and Future)

2. Senior, John M. 2003 *Optical Fibre Communications: Principles and Practice*, Prentice Hall, ISBN 0-13-635426-2
   A good and detailed introductory text for degree level study.

Assembly

1. http://om.pennnet.com/home.cfm
   A news and analysis site on optical assembly

2. http://www.circuitsassembly.com
   An electronics assembly web site

3. http://www.newport.com/store/product.asp?id=2839&lang=1
   A tutorial on optical coupling

Standards

1. Cunningham *et al.* 1998 *Gigabit Ethernet Networking* Macmillan ISBN 157800620
   A text book on the technical aspects of the GbE specification and the work involved to reach agreement on the standard.

2. 802.3AE-2002 IEEE Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications-Media Access Control (MAC) Parameters, Physical Layer and Management Parameters for 10 Gb/s Operation 2002/, ISBN 0-7381-3287-X IEEE Product No. SH94996-TBR IEEE Standard No.:802.3AE-2002
   The new 10 Gb/s Ethernet standard

Future, chip to chip

1. Savage, Neil 2002'Linking with Light', IEEE Spectrum, Vol. 39, No. 8, August 2002
   A magazine article describing the potential application of optical communication for chip to chip interconnection

2. http://www.spectrum.ieee.org/WEBONLY/publicfeature/aug02/opti.html
   A similar article available on the web

3. http://icd.pennnet.com/Articles/Article_Display.cfm?Section=Articles&Subsection=Display&ARTICLE_ID=123476
   An online article addressing the electrical and optical options for expanding the backplane limits

Future fibre to the home

1. Various authors 2001 *Topics in Broadband Access*, *IEEE communications magazine*, Vol. 39, No. 12, December 2001, pp. 116–134
Invited technical papers by industry experts on FSAN and related fibre to the customer initiatives

Glossary

10GbE the latest Ethernet specification, agreed in 2002 (purchase details are included above in Further reading

ADSL asymmetric digital subscription line, a modulation technique that allows extremely high amounts of data to be transferred using standard copper phone lines. Data rates range from 64 kb/s to 45 Mb/s at a fraction of the price of standard technology for the same speed. The technology is not widely available yet, since some countries are only in the early stages of rollout. There are many variants of DSL, called 'xDSL'

CATV cable television, which can be distributed by optical fibre. The digital data rate that can be supported by the CATV distribution network depends on the detailed design and the time it was installed, with more recent designs offering broadband services, not telephone only.

CWDM coarse wavelength division multiplexing, a method of combining light of different wavelengths to carry more data along the same fibre. Identical to WWDM, the channel spacing is wide to allow natural temperature variations

dB a ratio of a parameter, often power (3 dB indicates the power is halved)

dBm the power relative to 1 milliwatt, so 0 dBm is 1 mW

DFB distributed feedback laser, which has a grating included in the structure to define the operating wavelength more exactly and allow longer distance operation

DWDM dense wavelength division multiplexing, a method of combining light of different wavelengths to carry more data along the same fibre. The term 'dense' indicate that the wavelengths are close together giving more transmission capacity (and requiring thermal control)

ECOC European Conference Optical Communications, the major optical-communication conference, where the latest research results of workers in the field are presented, held annually in September

EFM Ethernet in the first mile, an IEEE standard under development, for the first mile from the customer, intended to provide fibre to the home

FSAN Full Service Access Network, is a standard that was started by telecommunications services providers to provide lower-cost fibre connections to the customer

Gb/s a data rate of 1000 million bits per second

Gbps a data rate of 1000 million bits per second

GbE Gigabit Ethernet a standard of the IEEE for 1-Gb/s data transmission over fibre using the principles of the lower-data-rate Ethernet used in most PC connections

IEEE (I-triple-E), a non-profit, technical professional association of more than 377,000 individual members in 150 countries, (but the USA base ensures a high proportion of American members). The full name is the Institute of Electrical and Electronics Engineers, Inc.

ISDN integrated services digital network, a digital service available over standard phone lines with two lines each of 64 kb/s. It has been available for some years but it has always been more expensive than standard phone lines (56kb/s is available from the latest modems for standard lines).

IST Information Society Technologies Programme, a single, integrated EU research programme building on the convergence of information processing, communications and media technologies

ITU International Telecommunications Union, the international standards body for the telecommunications industry

kb/s a data rate of 1000 bits per second

kbps a data rate of 1000 bits per second

LAN local area network, the data transmission hardware and software to distribute information locally, frequently around a building or small campus or several buildings of a business, often using only equipment owned by the business

Mb/s a data rate of 1 million bits per second

Mbps a data rate of 1 million bits per second

MSA multi-source agreement, a definition between competing industrial concerns to supply components with common outlines and interfaces, to make them mutually interchangeable

OIDA Optoelectronics Industry Development Association, a trade organisation body

OIF Optical Internetworking Forum, a body, mostly industrial, which supports the IEEE standardisation process by allowing detailed discussions amongst potential competitors and allowing consensus to be reached outside the formal process

PON passive optical network, a fibre distribution method where the signal is split between several users giving cost savings on the source, since it is shared between several users

Sematech the international organisation for the silicon industry, which ensures all parts of the equipment, processes, modelling tools etc. are in place for the next generation of integrated circuits

TDM time-division multiplexing, the method of combining electronic messages together from parallel electronic signal into a higher-speed serial signal.

Telcordia the short-form name for a set of reliability requirements and design specifications for use on the telecommunications network in the USA. It has become the de facto standard for reliability of optical components across the industry

VCSEL vertical-cavity surface-emitting laser, a laser from which the light output is emitted from the surface of the chip, with advantages of on-wafer test

WAN wide area network, the data transmission hardware and software to distribute information beyond a campus or several buildings of a business, often using the leased facilities of a telecommunications service provider such as BT

WDM wavelength-division multiplexing, a method of combining light of different wavelengths to carry more data along the same fibre

WWDM wide wavelength-division multiplexing, a method of combining light of different wavelengths, to carry more data along the same fibre; identical to CWDM, the channel spacing is wide to allow natural temperature variations

7.0 References

[1] Personick, Stewart D. 2000 'Evolving toward the next generation Internet, Challenges in the Path Forward', IEEE Communications Magazine, Vol. 40, No. 7, July 2002 (Special Issue on The Internet Past Present and Future)
[2] Optical Fibre Communications Principles and Practice (see Further reading, p. 36)
[3] http://www.ntt-at.com/products_e/mtconnect/
[4] 'High Capacity Polymer Optical Fibre Systems' Khoe et al ECOC 02 paper3.4.1 invited review
[5] http://www.plastics-car.com/applications/electrical.html
[6] 10Gb/s Uncooled Laser diodes for Datacommunication Application Meliga et al Agilent ECOC 02 paper 5.3.1
[7] Subcarrier multiplexed lightwave systems for broad-band distribution
GTE Lab. Inc., Lightwave Technology, Journal of, p1329–1342, Sept. 1989, Vol. 7, No. 9
[8] http://www.paralleloptics.org/
[9] Demonstration of a small-form-factor WWDM transceiver module for 10-Gb/s local area networks
Buckman, L.A.; et al IEEE Photonics Technology Letters , Vol. 14, No. 5 , May 2002, p 702–704
[10] http://www.finisar.com/media/product_document_detail/site2_902409951_FWSMUX-DEMUX.pdf
[11] http://www.broadcom.com
[12] http://www.velocium.com/templates/cusvelocium/default.asp?id=22211
[13] http://www.opnext.com/
[14] http://www.gtran.com/
[15] http://www.gcsincorp.com/
[16] http://www.inphi-corp.com/
[17] http://www.infineon.com/cmc_upload/documents/037/543/_V23814-15_U1306_M130_011128.pdf
[18] http://www.semiconductor.agilent.com/cgi-bin/morpheus/promotions/promotionsDetail.jsp?promoID=20163
[19] Parallel optical interconnections development at HP Labs
Yuen, A. et al, Lasers and Electro-Optics Society Annual Meeting, 1997. LEOS '97 10th Annual Meeting. Conference Proceedings., IEEE p 191–192, Vol. 2
[20] Module packaging for high-speed serial and parallel transmission
Karstensen, H. et al Infineon Technol. AG., Berlin, Germany Proc ECTC 2000. :p 479 - 486 21-24 May 2000
[21] http://www.newport.com/file_store/PDFs/tempPDFs/LaserWeld_e4438.pdf
[22] Agilent technology's single-mode small-form-factor (SFF) module incorporates micromachined silicon, automated passive alignment and non-hermetic packaging to enable the next generation of low-cost fibre-optic transceivers
Owen, M. Advanced Packaging, IEEE Transactions on, Vol. 23, No. 2 , May 2000, pp. 182–187
[23] http://www.loctite.se/tds/EO1061.pdf
[24] Compound Semiconductor 'Low cost manufacturing is key to the future of photonics', vol. 8, no7, Aug. 2002, p. 37
[25] III-Vs review 'Photonics enter the automation era' Vol .15, No. 9, Nov/Dec 2002, p. 42
[26] http://www.tycoelectronics.com/prodnews.asp?id=364
[27] http://www.eetuk.com/bus/news/OEG20021112S0003
[28] http://www.freedoniagroup.com/browse.html report #1550 06/2002
[29] Cunningham *et al.* 1998 *Gigabit Ethernet Networking*
[30] http://standards.ieee.org/getieee802/802.3.html
[31] http://www.xenpak.org/MSA/XENPAK_MSA_R2.1.pdf
[32] http://www.xfpmsa.org/cgi-bin/home.cgi
[33] http://www.xpak.org
[34] http://www.x2msa.org/
[35] http://www.blazenp.com/news/pr_12_18_01.html
[36] ECOC 02 paper 8.2.1 Sogawa et al Sumitomo
[37] http://www.oiforum.com

[38] http://www.fsanet/
[39] http://grouper.ieee.org/groups/802/3/efm/baseline/opticsbaseline.html
[40] '1.02 Tb/s, 3km Multi-mode fibre Feasibility-Experiment using 200-channel DWDM and Quadrature-Subcarrier Transmission' ECOC 02 paper 8_2_2
[41] IEEE Communications Magazine, Vol. 40, No. 7, July 2002
[42] http://www.cordis.lu/ist/home.html
[43] http://www.rci-france.com/
[44] http://public.itrs.net/Files/2003ITRS/Interconnect2003.pdf, p. 37
[45] http://public.itrs.net/Files/2003ITRS/Interconnect2003.pdf, p. 15
[46] http://dbs.cordis.lu/fep-cgi/srchidadb?ACTION=D&CALLER=PROJ_IST&QF_EP_RPG=IST-2000-28358
[47] http://dbs.cordis.lu/fep-cgi/srchidadb?ACTION=D&CALLER=PROJ_IST&QF_EP_RPG=IST-2001-35235
[48] http://www.intec.rug.ac.be/io/publications.html
[49] http://www.phy.hw.ac.uk/resrev/star/index.html
[50] http://www.e2oinc.com/, http://www.bigbearnetworks.com/. Both companies have headquarters in California
[51] http://www.optillion.com/
[52] http://www.teraconnect.com/product.shtml
[53] http://www.blazenp.com/
[54] http://www.vertilas.com/
[55] http://www.broadcom.com
[56] http://www.phyworks-ic.com/
[57] http://www.digitaloptics.com/
[58] http://www.opticalantennasolutions.com/