

This is Chapter 4 from the second edition of :

Networks, Routers and Transputers: Function, Performance and applications

Edited by M.D. May, P.W. Thompson, and P.H. Welch

© INMOS Limited 1993

This edition has been made available electronically so that it may be freely copied and distributed. Permission to modify the text or to use excerpts must be obtained from INMOS Limited. Copies of this edition may not be sold. A hardbound book edition may be obtained from IOS Press:

IOS Press
Van Diemenstraat 94
1013 CN Amsterdam
Netherlands

IOS Press, Inc.
P.O. Box 10558
Burke, VA 22009-0558
U.S.A.

IOS Press/Lavis Marketing
73 Lime Walk
Headington
Oxford OX3 7AD
England

Kaigai Publications, Ltd.
21 Kanda Tsukasa-Cho 2-Chome
Chiyoda-Ka
Tokyo 101
Japan

This chapter was written by H. Gurney and C.P.H. Walker.

4 Connecting DS-Links

4.1 Introduction

Digital design engineers are accustomed to signals that behave as ones and zeros, although they have to be careful about dissipation and ground inductance, which become increasingly important as speeds increase. Communications engineers, on the other hand, are accustomed to disappearing signals. They design modems that send 19200 bits per second down telephone wires that were designed 90 years ago to carry 3.4KHz voice signals. Their signals go thousands of kilometers. They are used to multiplexing lots of slow signals down a single fast channel. They use repeaters, powered by the signal wires.

Digital designers do not need all these communications techniques yet. But sending 100Mbits/s or more down a cable much longer than a meter has implications that are more analog than digital, which must be taken care of just like the dissipation and ground inductance problems, to ensure that signals still behave as ones and zeros.

Actually, it is easy to overestimate the problems of these signal speeds. Engineers designing with ECL, even fifteen years ago, had to deal with some of the problems of transmitting such signals reliably, at least through printed circuit boards (PCBs), backplanes, and short cables. One of the best books on the subject is the Motorola 'MECL System Design Handbook' [1] by William R Blood, Jr., which explains about transmission lines in PCBs and cables. This shows waveforms of a 50MHz signal at the end of 50ft (15m) of twisted pair, and of a 350MHz signal at the end of 10ft (3m) of twisted pair, both with respectable signals.

This chapter first discusses the signal properties of DS-Links. PCB and cable connections are then described, followed by a section on error rates: errors are much less frequent on transputer links than is normal in communications. A longer section introduces some of the characteristics of optical connections including optical fibre, which should be suitable for link connections up to 500m, using an interface chip to convert between the link and the fibre. A pointer is given towards possible standards for link connections. Appendix A describes a connector that will assist standardization of transputer link connections. Appendix B shows waveforms of signals transmitted through cable and fibre. Appendix C gives detailed electrical parameters of DS-Links, and appendix D gives an equivalent circuit for the DS-Link output pads.

4.2 Signal properties of transputer links

Considerable design work has gone into making the DS-Link signals [4] well behaved. The bit-level protocol and the electrical characteristics both contribute to make the link signals unusually easy to use, for serial data at 100Mbits/s.

The DS-Link information is carried by a pair of wires in each direction. The D signal carries data bits, and the S signal is a strobe, which changes level every bit time that the D signal does not change¹². This is illustrated in figure 4.1. This bit-level protocol guarantees that there is a transition on either D or S every bit time. Effectively this provides a Gray code between the D and S signals.

12. Note that this differs from the usual meaning of a 'strobe', which is a signal which indicates *every* time the data signal is valid.

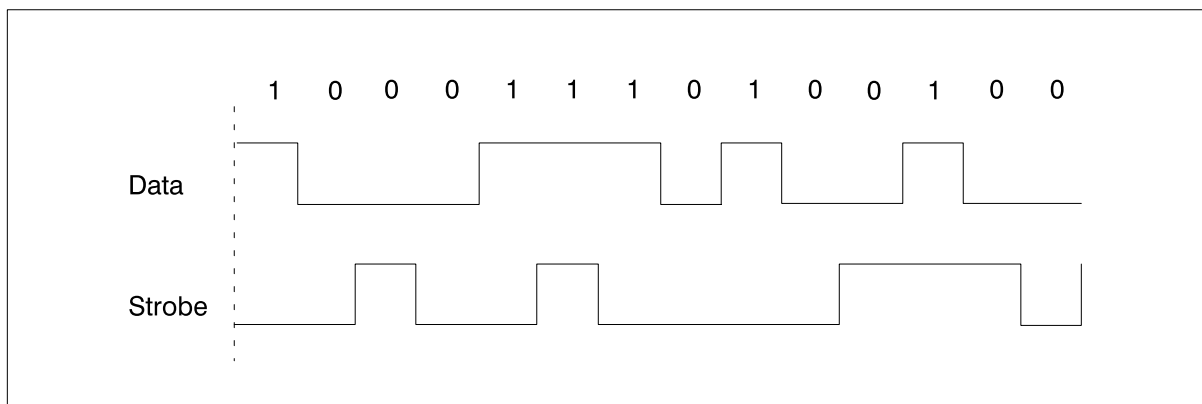


Figure 4.1 DS-Link signals

One result of the DS Gray coding is that the received data is decoded purely from the *sequence* of D and S transitions rather than depending on any absolute time. This means that the link receivers ‘autobaud’, receiving data at whatever speed it is sent (so long as the receiver logic is fast enough).

The Gray coding makes it much easier to design logic that is fast enough, because the timing resolution required is a whole bit time. Alternative codings would require a clock edge in the centre of a data bit, and hence require timing resolution of half a bit time. The more relaxed timing resolution needed by the DS-Links gives major benefits in terms of the performance that can be achieved in practical systems.

A further advantage of the coding, with only D or S changing at a time, is that the signal can be received without a phase-locked loop (PLL) – the clock is just the Exclusive-OR of the D and S signals. For the C104 routing switch, avoiding the need for 32 PLLs is very valuable, and it is likely that a 32 way routing switch would not be implementable had the PLLs been required.

Electrical aspects of the design include a controlled output impedance approximately matched to a 100Ω transmission line¹³. Obviously there is a tolerance on the impedance, which also may not be identical for high and low, but the DS-Link has been designed to minimize the effect of any such mismatch on the signal.

The link outputs have also been designed to give controlled rise and fall times. The full electrical characteristics will not be known until the devices are fully characterized, but a reasonable estimate of the transition times is 3ns fastest transition and 6ns slowest transition.

The DS coding gives as much tolerance as possible for skew between the D and S signals, and the outputs and inputs have been designed to have minimal skew at the TTL threshold of 1.5V.

These characteristics of the DS-Link signals make them ideal for connections on PCBs, and for DC coupled connections on short lengths of cable, up to 10m. Later sections will describe such connections, as well as much longer connections up to 500m using optical methods.

4.3 PCB connections

The following discussion assumes the use of multi-layer PCBs with power and ground planes; use of DS-Links on double-sided boards without ground planes is not recommended.

A 100Ω transmission line impedance is fairly easy to achieve on the surface of a PCB. PCBs have been made with long connections of 100Ω impedance which carry link signals faithfully. The 100Ω impedance requires a track width between 0.1mm and 0.3mm, depending on the board

13. See appendix C.

thickness and where the power planes are located within it. Figure 4.2 (derived from data given in Blood [1], from SONY [2], and from Coombs [3]) shows the approximate relationship between these parameters for standard FR4 PCB material with a dielectric constant of 4.7.




Figure 4.2 Graph showing approximate PCB transmission line impedance for FR4 laminate

Note that when a PCB track is buried in the fiberglass/epoxy laminate, its impedance is reduced by about 20% compared with a surface track. This requires the inner layer tracks to be narrower than surface tracks, to minimize differences in impedance. It is not possible, within the normal 1.6mm board thickness, to have 100Ω tracks sandwiched between power or ground planes.

If the transmission line impedance could be maintained with high precision, PCB DS-Link connections would be good for several meters, in theory. However in practice it is hard to maintain a tighter tolerance than $\pm 20\%$. It is therefore advisable to limit the connections on PCBs to less than 1000mm with standard FR4 PCB material. If the impedance goes outside the range of 80Ω to 120Ω, it is advisable to limit the connection to 500mm.

Short discontinuities in the impedance are permissible, such as connectors, vias, and short sections of track of higher or lower impedance; such discontinuities should be kept to less than 50mm. Similarly, if it is necessary to use some PCB tracks of higher impedance than 100Ω, and some lower than 100Ω, it is best if they can be alternated in short sections, rather than having a 400mm length of 125Ω track and then a 400mm length of 80Ω track.

The controlled transition times of the DS-Links minimize crosstalk compared with the sub-nano-second fall times of some of the fast families of 'TTL', but care still needs to be taken over crosstalk. Tests, simulations, and theory using typical PCB materials and DS-Link outputs suggest that backwards crosstalk increases as the length of the parallel tracks increase up to 25cm, and

does not increase for longer parallel tracks. Track separation of 0.15mm over this length appears to give 1 volt of crosstalk, which is above the noise margin. Simulations of track separation of 1.25mm over a length of 20cm give crosstalk figures of less than 100mV.

The references [1], [2], and [3] do not give a great deal of information about PCB crosstalk, and the results above suggest that further work is required. In the meantime, it must be good practice to avoid long parallel runs and to space the tracks out as far as possible. Another technique is to use guard tracks (tracks connected to 0V) between link tracks, although the effects of this on the impedance of the link track may need to be taken into account.

The D and S pair of signals should be approximately the same length, but a difference in length of 50mm would only introduce a skew of 250ps, which should be totally acceptable.

4.4 Cable connections

This section looks at existing cable interfaces, comparing them with transputer links, and then discusses the loss and noise that occur in a cable, and what can be done to overcome their effects.

4.4.1 Existing cable interfaces and rough costs

Ethernet connections are now inexpensive, with a component cost well under \$50 and an end-user cost around \$150. Transputer links are even less expensive with a low cost T400 having two OS-Links each capable of 20Mbits/s full duplex, a total bandwidth four times that of ethernet.

Token Ring goes a little faster than Ethernet, but to go substantially faster the next standard is FDDI at 125 Mbits/s (of which 100 Mbits/s are useful data). FDDI is expensive, not only in its protocol, but even in its components, and just the optical transceiver is not expected to fall below \$100 even in volume for some time.

Links on the T9000 transputer run at 100 Mbits/s, full duplex. The cost per link is considerably less than either the chipset or the transceiver for FDDI. The C104 routing switch, with 32 ports will give a cost per port well under \$10 – at least an order of magnitude less than the FDDI component cost.

Ethernet, Token Ring, and FDDI are all local area networks, with many ports in a network and long distances between ports. Transputer links are point-to-point, and are generally expected to be comparatively short connections. In this respect they are more like the recent parallel interfaces such as SCSI2, IPI and HPPI. HPPI as an example has a maximum length of 25m, and runs at 800 Mbits/s in one direction down a cable with 50 twisted pairs. The same speed in both directions requires two cables, and the speed can be doubled by using two cables in each direction.

FibreChannel is a fibre connection with similar data rates to HPPI, using laser diodes. This will allow much longer connections than HPPI, at drastically lower cable costs, but possibly with a high cost per port.

4.4.2 Vanishing signals (High frequency attenuation)

Copper wire has a finite resistance: 28AWG wire is one of the smallest cross sections in widespread use and has a resistance of $0.23\Omega/\text{m}$, 1Ω in 4.3m. If the characteristic impedance of the cable is 100Ω , a resistance of 10 ohms is not going to affect the signal very much, so this cable should certainly be usable at 43m. The problem is that at high frequencies, the signal does not flow evenly throughout the conductor but concentrates at the outside of the conductor – the skin effect. So the higher the frequency of the signal, the more the resistance of the cable. Some of the energy does not flow in the conductor at all, but in the insulation and, if it can, in adjacent

conductors causing crosstalk as well as loss. Some of the energy is sent into the atmosphere to interfere with radios and other users of the airwaves.




Figure 4.3 Cable attenuation against frequency for a variety of cables

The sum of these losses of energy which depend on frequency is measured in dB (deciBels) per unit length. Figure 4.3 plots these losses for a number of cables – some inexpensive, some extremely expensive. The detail is not important but note that for all of the electrical cables the loss increases with frequency.

The increase of loss with frequency means that the higher the frequency to be passed along the cable and the longer the cable, the less ‘lossy’ (and probably more expensive) the cable will have to be. Above some length of connection, the losses have to be compensated for somehow – as in Telecommunications – and more tricks have to be used, increasing the cost of the circuitry at the ends of the cable, and possibly adding repeaters in the cable. At some stage, it will become worth while to use optical fibre, an example of which is shown in figure 4.3.

The increased loss at high frequency can be overcome by using a cable short enough that the loss is minimal. At 100MHz, this could mean less than a meter for some of the cables illustrated. The effect of using a longer cable is distortion of the signal. Figure 4.4 shows the sort of thing that happens to an NRZ (Non Return to Zero) signal which has suffered a 10dB loss¹⁴ at the frequency of the square wave. The dotted line represents the DC threshold of the receiver, which suggests that the signal will not be received correctly, even if there is no noise.




Figure 4.4 Cable as a low-pass filter

Figure 4.5 shows a similar effect to figure 4.4, but the received high frequency voltage is now about 0.6 times the transmitted voltage, representing a loss at this frequency of around 4.5dB. At 100Mbps/s, the ‘sine wave’ part of figure 4.5 is 50MHz, and the 28AWG IPI/SCSI2 cable¹⁵ shown in figure 4.3 has a loss of 2.8dB for 10m at 50MHz, so in the absence of noise, and with a receiver which had sufficient gain and is tolerant of small errors in timing, this cable might just work not at 43m but at $(4.5/2.8) \times 10\text{m}$ or 16m. In practice the maximum length will be less than this.

14. The 10dB loss means that the power at the receiver is 1/10 of the power at the transmitter. As power is volts times amps, both of which are reduced in the same proportion, the received voltage for a 10dB loss is 0.33 times the transmitted voltage.

15. The 28AWG and 26AWG Madison cables, shown in figure 4.3, have also been designed to minimize the skew that can occur between any two pairs in the cable, resulting in a skew of 0.04ns/ft, which is an order of magnitude better than that of cables which have not been so designed. Skew is important for parallel interfaces such as SCSI or HIPPI, and is equally important for DS-Links.




Figure 4.5 Almost enough signal

4.4.3 Boxes are not at the same voltage (Common mode, DC coupled differential signals)

For a cable several meters long between two boxes, there may be crosstalk and it can not be guaranteed that there will be no difference between the boxes' ground or logic 0V levels. Any difference will be seen as noise.

A good way to remove the effect of the difference in grounds between the two boxes is to send differential signals. These are shown in figure 4.6. Any difference in ground voltage will be seen as common mode by a receiving differential buffer.




Figure 4.6 DC coupled differential signal

A popular standard differential signals is RS422, whose receivers have a common mode tolerance of $\pm 7V$. The RS422 components are limited to 10Mbits/s or 20Mbits/s, and so are not suitable for higher bit rate DS-Links. However they have been found to be extremely reliable when used to connect OS-Links between boxes, which shows that differential signalling is effective. DS-Links therefore simply require faster differential buffers.

ECL buffers are much faster than the RS422 components. Blood shows 'scope traces of a 350MHz signal after a receiver at the end of 10ft of twisted pair. Unfortunately the ECL common mode tolerance is much less than RS422, from +1V to -1.8V or -2.5V depending on the device used.

A family of devices from AT&T (41 series of High Performance Line Drivers, Receivers, Transceivers.) offers speed approaching that of ECL together with common mode tolerance approaching that of RS422. The transmitters have TTL inputs and pseudo-ECL outputs, and the receivers convert the pseudo-ECL back to TTL. One range of devices runs up to 100MHz (200Mbits/s), another to 200MHz (400Mbits/s). Common mode tolerance is from -1.2V to +7.2V, with the 1V signal approximately in the middle of of this range.

Tests have been done using these buffers which indicate that a 10m link running at 100 Mbits/s should work reliably.

The cable used for the tests was 30AWG individually shielded twisted pairs. The shielding and the use of 30AWG both increase attenuation compared with the 28AWG unshielded cable mentioned earlier; the shielding minimizes EMC emissions for FCC and other regulations, eliminates crosstalk, and the 30AWG reduces the size of the cable.

4.4.4 Ground differences more than a few volts (AC coupling, DC balance)

In the last section, we overcame some problems by using balanced, differential signals. Larger common mode voltages between two boxes can be overcome by using AC coupling, which requires a different sort of balance. Figure 4.7 shows a signal which has a mark-to-space ratio of 4-1: on the receive side of the AC coupling, the threshold is set by averaging the received voltage. As a result, the threshold is heavily offset, reducing the noise margin and changing timings.




Figure 4.7 Effect of DC imbalance

In order to provide DC balance, so that the threshold is in the middle of the signal, the data is *coded* in some way, usually by adding redundant bits to achieve the desired signal characteristics. One of the most popular forms of DC balanced coding is Manchester Code, which provides DC balance over every bit period, at the expense of doubling the bit-rate. An alternative to coding is to *modulate* a carrier, in amplitude, in frequency, in phase, or in combinations of these, with different data values being represented by different amplitudes, frequencies or phases; the carrier is a sine wave which is inherently DC balanced.

Even when there is no DC component in the signal, a long period without a transition can cause the signal to disappear. Codes therefore have a maximum *run length* to limit this time between transitions; they also have a minimum run length, to ensure that two adjacent edges do not cancel each other out and appear as no edge. Figure 4.8 shows the effect of a long run length: the signal droops, reducing the margin between the signal and the threshold, until it eventually crosses over the threshold.




Figure 4.8 Effect of excessive run length

Some of the codes which are currently popular are not in fact completely DC balanced, but for most data patterns have minimal DC component. Such codes include the 2:7 *Run Length Limited* code used on disks, and the TAXI/FDDI code which is never worse than 40%/60% balanced. (The code used by FibreChannel has the same efficiency as the FDDI code, but is completely DC balanced.) A technique used on ISDN and on SONET is to *scramble* the data so that it is approximately balanced and very rarely has long run lengths; the scrambling has the advantage that no extra bits are added to the data.

An extreme form of AC coupling is to differentiate the signal, which provides inherent DC balance. The pre-compensation circuit used in twisted pair FDDI effectively produces the sum of the signal itself, plus a differentiated version of the signal. In magnetic recording, such differentiation occurs naturally, but it brings its own problems; any noise such as crosstalk is coupled through the differentiator, and any AC imbalance in common-mode coupling is translated into extra noise.

AC coupling can either be provided by transformers or capacitors. Transformers provide excellent common mode isolation and are readily available at low cost up to a few hundred MHz (Mini-Circuits T1-1 0.15MHz to 400MHz, \$3.25 in low volume). Capacitors do not provide good common mode isolation, but can be used for frequencies up to many GHz. Low cost amplifiers are also available which *must* be AC coupled, with 7.5dB gain at 1.5GHz.

4.4.5 Limiting the frequency range and tuning

The constraints on run length and on the DC balance effectively reduce the bandwidth that needs to be received. If the highest frequency needed is 50MHz, and the lowest is 10MHz, the 28AWG cable referred to above loses 2.8dB in 10m at 50MHz and 1.2dB at 10MHz. So we only have to cope with a difference of 1.6dB per 10m between the frequencies. Instead of the 16m limit given above for the differential DC coupled case giving 4.5dB, we can AC couple, use more gain, and should be able to reach $(4.5/1.6) \times 10\text{m}$ or 28m¹⁶.

Even with a very wide bandwidth, it is possible to use tuning to compensate for the frequency characteristics of the cable. As with 'scope probes, it is easier to do if the tuning is built into the cable (otherwise it has to cope with a wide range of different cable lengths). As with 'scope probes, this can be expensive and liable to misuse.

16. A similar example is Ethernet, which uses Manchester coding, with a limited frequency range, and allows a total of 8.5dB loss at its frequency.

4.5 Error Rates

The form of serial communications that most engineers are familiar with are LANs and very long distance (tele-)communications. For these long distance connections, error rates tend to be around 10^{-9} or less, which at 100MBits/s is an error per link every five seconds (counting a link as bidirectional). Telecomms and LANs also need to cope with buffer overflow.

For these high error rates, it is absolutely necessary to have CRCs for error detection, and to have re-try mechanisms for corrupted or lost data – whether lost as a result of data errors or buffer overflow.

Another reason for needing CRCs is that most of the efficient communication codes, such as FDDI and FibreChannel, allow an erroneous single bit in the received data stream to be decoded as a valid (but incorrect) whole data symbol; both the FDDI and FibreChannel codes limit such decoded errors to less than a byte of data, but such error multiplication necessitates the use of checksums such as CRC.

The situation with transputer links is rather different: the specified error rates on PCBs are substantially better than 10^{-20} , which is a failure per link every 50 000 years. At such error rates, it is quite reasonable to consider a system as reliable, and to crash the system if an error occurs. Alternatively, it is possible to add software to detect the rare event and to take some form of recovery action. In practice, at these error rates, hardware errors are much more likely to be caused by lightning strikes or by mechanical damage than by electrical signal failure.

The parity check on the DS-Links is such that a single bit error, either in control or data, is detected. As long as the errors are infrequent (one every several thousand years), this is entirely adequate. If a user is concerned about the possibility of an error not being detected, software can be added to the processes at the end of the link to perform more rigorous data checks and to recover from data or control errors.

These software checks can be performed even if the suspected virtual channel goes through a routing switch. The suspected link can be configured in the routing switch to go to a single transputer which is programmed to check the messages, effectively ignoring a possibly corrupted routing header. If several transputers are programmed to check the messages, the routing switch can be configured to route the messages to any of these transputers – but not to another routing switch or to a transputer that is unable to check the message.

The specifications stated in the transputer data sheets are designed to ensure the very good error rates that are expected between logic devices on a PCB. As a result, the permitted skew specification for the T4xx and T8xx transputers is a few nanoseconds. Some users have observed that OS-Links work with larger skews, but with such large skews the error rates are more like the 10^{-9} of the telecommunications and LANs. At INMOS, there is a network of transputer links, buffered with RS422 buffers, with connection lengths of close to 100m – far outside the specification or recommendations; in practice, the incidence of software failure on this network is substantially higher than the incidence of hardware errors due to links.

DS-Links have been specified, therefore, so that they give such infrequent errors that the hardware can be considered reliable. This does not preclude any user from adding checking software; nor does it preclude the use of more elaborate checking hardware when connecting links over longer distances such as with optical fibre interconnections.

4.6 Optical interconnections

Included in this section on optical interconnection are optical isolators which retain electrical connection, but offer large tolerance of common mode noise, and optical fibre, which comes into its own for connections much above 10m.

4.6.1 DC coupling with common mode isolation (Optical Isolation)

Optical isolators appear to offer the best of both worlds, in that they do not require the DC balance or run length limits that AC coupling needs, but yet offer almost infinite tolerance to common mode. To make opto-isolators fast, however, most of the circuitry needs to be included that would be used in an optical fibre connection. As a fibre connection would cost less than the wire connection and go much further at a given speed, it may be preferable to use fibre. Whether this is the reason or not, it has not been possible to find opto-isolators that are specified to run at 100Mbits/s.

4.6.2 Long distance, high data rate, infinite isolation ... but... (Optical Fibre)

The fibre shown on figure 4.3 is inexpensive but is much better in terms of its attenuation than the best copper cable. Single mode fibre is still better. The problem is not in the attenuation in the cable, but in the losses (and consequent costs) in converting from electricity to light at one end and from light to electricity at the other end.

4.6.3 Losses, performance and costs of components for optical fibres

The light is produced by a LED or by a Laser Diode. An example LED outputs (infra-red at 1300nm wavelength) 0.25mW of optical power when driven by 100mA of electrical power. Laser diodes are more efficient, one for example produces 5mW of optical power for 50mA of input current. The fastest LEDs have an optical rise time of about 2.5ns, and a 1.5dB cutoff at 100 or 150MHz (6dB around 800MHz). The 1300nm laser diodes have sub-nanosecond rise and fall times: one example has a very sharp cutoff at around 1.5GHz.

Components with wavelengths of 820 or 850nm are in many respects more suitable for 100 Mbits/s transputer links. Components from HP and from a number of other companies include LEDs which output around 0.1mW (-10dBm) of optical power into the fibre with optical rise and fall times of 4ns, for a current of 60mA.

The receivers are PIN¹⁷ photodiodes, very often integrated into a hybrid with a pre-amp, and sometimes also with a power supply for the diode. The diodes are reverse biased, with a finite reverse (Dark) current. One example has a responsivity of about 0.5A/W. Assuming no attenuation in the fibre, 100mA into the LED becomes 0.25mW in the fibre which becomes 0.125mA in the PIN diode; this loss is far more than the electrical cable loss but fibre has the important advantage that, over short distances at least, there is much less variation of loss with frequency.

The received current needs to be amplified up to logic levels, and this amount of amplification, at these frequencies, is easier with AC coupling. So the requirements of bandwidth limiting, DC balance and run length limiting are present for optical fibre as much as for electrical wire. The FDDI transceivers and the HP 820 nm 125MHz receiver module amplify up the current into a voltage – ECL levels from the FDDI transceivers, 10mV to 1V from the HP receiver.

The costs are radically dependent on the technology used, as illustrated in table 4.1 (all figures are approximate and for large volumes).

17. PIN = P doped, Insulator, N doped

Table 4.1 Optical components cost/performance

Wavelength (nm)	Data rate	Light source	Cost	Availability
820	200KBits/s	LED	less than \$10 per transceiver	now
820	125MBits/s	LED	\$30 per transceiver	now
1300	125 to 350 MBits/s	LED	over \$300 per FDDI transceiver	now
			\$100 per FDDI transceiver	long term goal
1300	125MBits/s to 2.5GBits/s	Laser diode	\$1000 to \$10000 per transceiver	now

Notice that there is nearly an order of magnitude cost difference between the 820nm and 1300nm wavelengths, and another order of magnitude between LEDs and lasers. The one exception to this is the 780nm laser diodes used for Compact Disks, which are discussed below.

4.6.4 Expensive or affordable, long or short distances, 1300 or 820nm?

Most of the work on fibre has have been to make it go long distances, often at very high speed; or to make it cheap, where speed and distance do not matter. FDDI seems to come in between these, in asking for 2km at 125Mbits/s, but they have chosen the more expensive 1300nm. In fact FDDI connections using lasers are now being developed to go further than the 2km, as Medium or Metropolitan Area Networks (MANs).

The 820nm components are limited in distance to about 500m at 100 or 125 MBits/s, which is more than adequate for transputer links.

The laser diodes that are used in compact disks have a wavelength of 780nm, which ties in well with the HP 820nm receivers for 100Mbits/s, and it is possible that the CD lasers could be used with faster receivers to provide 400Mbits/s. FibreChannel has specified a CD laser as one of its options. These laser diodes are inexpensive because they are made in such large volumes for CDs, but the laser is not ideal for use by non-experts, and the laser diodes are not as reliable as LEDs.

At present, the cost, availability, and performance of the 820nm components appear to offer the preferred choice for DS-Links.

4.6.5 Interfacing between links and fibre

The last few subsections have described a number of characteristics of the fibre connection which are not handled directly by the DS-Link:

- The fibre connection is a single fibre in each direction, so both D and S need to be encoded onto a single signal;
- This signal needs to include sufficient transitions that a clock can be extracted by a PLL at the receiver;
- The LED (or laser) is driven by a current rather than by a voltage, and the receiver needs to see a signal of possibly only 10mV, certainly no more than ECL;
- The fibre allows connection up to 500m, whereas the buffering in the standard link logic is enough for some distance between 10m and 50m.

- Longer distance connections, with the amount of amplification required for the optical signal, is such that the connection must be considered as less reliable than normal short connections on a PCB. In fact the indications are that it may be possible to achieve worst case error rates of the order of 10^{-20} , far better than is achieved by normal communications. It may nevertheless be reasonable to offer additional error checking and possibly alternative means of handling errors compared with short distance links.

The best way to do these various interfacing functions would be with a link-to-fibre interface chip, designed for the purpose.

INMOS is collaborating on projects in the European ESPRIT program with other partners developing optical fibre connections. Indications suggest that fibre connection over 200m to 500m will be achievable with low-cost optical components. The signalling system used for the optical connection should allow isolated copper connection over 100m, possibly with unshielded twisted pair cable.

4.7 Standards

A number of users have asked that standards for interconnections between equipments be proposed, so that different manufacturers' equipments can be connected by their transputer links. In some respects this provides a 'small area network' of transputer or link based systems.

The proposal for electrical cable connection is to use DC coupling with the 41 series buffers mentioned earlier. Earlier in this chapter, it was suggested that these cable connections should work well up to 16m, and although tests have given good results at 30m, for a reliable link it is necessary to limit this to 10m using the 30 AWG shielded twisted pair cable suggested.

If isolation is required the proposal is that it should be done with low cost optical fibre.

In drafting early versions of the proposed standard, it was found to be necessary to specify four different types of connector for different applications. There was no single connector which provided separate cables for each link, while meeting the other requirements, so INMOS produced an outline specification of a single connector which would satisfy all the various requirements. This connector has been developed by AMP, Harting and Fujitsu, in cooperation with INMOS/SGS-Thomson. Plugs and intermateable sockets have been manufactured by Fujitsu and Harting, and the connector closely follows an IEC standard which was originally put forward by AMP. It is shielded, polarized, latched and robust, and has a leading pin for 0V for reliable hot-swap. An outline description of this connector is included as an appendix.

The four connectors specified in the draft standard were 9-way D type, LEMO, SCSI2, and METRAL. Pinouts will be defined for these, for the MiniDIN, and for the new connector.

Proposed standards for optical fibre connection are based on a fibre interface chip, with the low cost 820nm optical components, 62.5 μ m fibre (which is being installed into buildings for FDDI) and SC connectors (which appear to give a good combination of repeatability, density, and ease of use for the end-user).

The electrical and optical issues covered by this chapter, the protocols of Chapter 3, and the connector of Appendix A are combined in a draft IEEE standard, P1355.

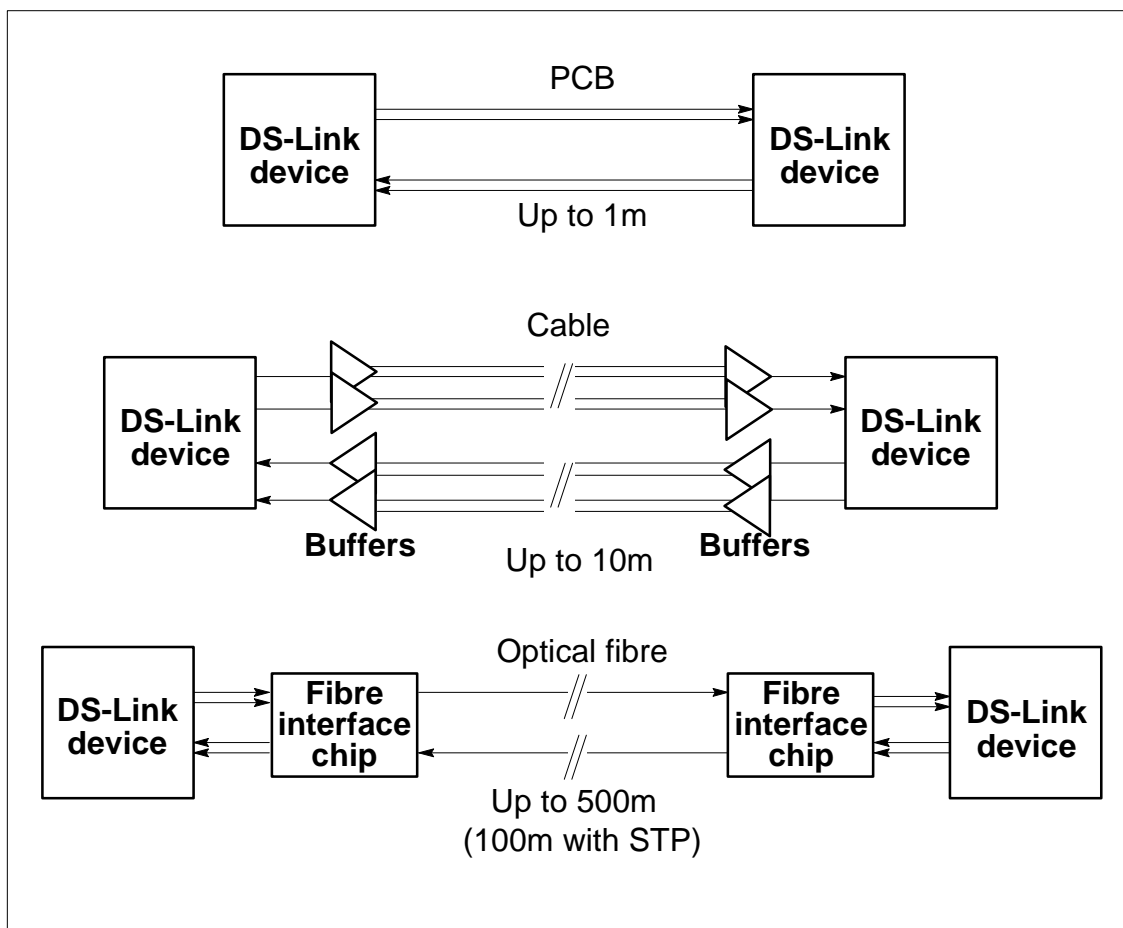


Figure 4.9 Distances that can be covered by DS-Links

4.8 Conclusions

DS-Links have been optimized for short connections on printed circuit boards, for which they are ideal. The Gray coding means that the receiver does not need a PLL, that there is a wide tolerance of skew, and that the receivers can ‘autobaud’ without requiring a status register to set their speed. The comparatively slow edges – at least for 100 Mbits/s – minimize crosstalk.

Link specifications are designed to ensure that errors are sufficiently infrequent that connections can be treated as logic connections rather than as telecommunications or LAN connections. If users violate these specifications for links, systems will often work, but with error rates approaching the error rates seen by LANs. For these error rates, it is necessary to add software to handle the more frequent errors. Such software is not required when the specifications are met.

For PCB connections up to 20cm, the characteristic impedance of the PCB track is not critical. Up to 1m the impedance should be kept within a reasonable tolerance, between 80Ω and 120Ω. Some care should be taken to avoid crosstalk. Beyond 1m, PCB connections may be possible, but the characteristic impedance should be more tightly controlled.

INMOS will be proposing link standards for long distance connections. Such standards will enable different manufacturers’ equipments to interconnect and, with cooperation on software, to inter-operate.

The proposal for short cable connections up to 10m is to use the fast 41-series buffers from AT&T, which have good common mode performance, in a DC coupled arrangement. For longer connections, up to 200 or 500m, or for electrical isolation, it seems best to use low cost optical fibre components, with a purpose designed interface chip.

Standards remove from the user some of the need to understand fully the principles on which they are based. At 100 Mbits/s, over the distances suggested here, the problems are not especially severe, but the faster the signals and systems go, the more necessary it is to engineer them to avoid problems such as attenuation in the connection. It is hoped that this chapter is of assistance in understanding these issues.

4.9 References

- 1 *MECL System Design Handbook*, William R Blood, Jr, Motorola.
This is an excellent book on the subject of high frequency digital logic signals on PCBs and cables. It also shows that the ECL system builders needed careful thermal design some years ago.
- 2 *SONY data book of SPECL*, 1990 edition.
This has a short application note with some comprehensive graphs of transmission line impedance, capacitance, and delay.
- 3 *Printed Circuit Handbook*, third edition, edited by Clyde F Coombs, Jr, McGraw-Hill, New York, 1988 ISBN 0-07-012609-7.
This book covers all aspects of printed circuits.
- 4 *The T9000 Transputer Products Overview Manual*, INMOS/SGS-THOMSON, 1991, order code DBTRANSPST/1.

There are many textbooks on communications but one of the most useful, which explains the concepts for a non-specialist and without excessive mathematics, is the Open University course 'T322: Digital Telecommunications'; this comprises a number of books, which are available separately or as a set from Open University Educational Enterprises in Milton Keynes, England. The three most useful in the course are Blocks 4, 5, and 6: Digital Signals; Noise; Coding and Modulation.

More mathematical, and covering more ground, is 'Digital Communication' by Edward A Lee and David G Messerschmitt, ISBN 0-89838-295-5, reprinted 1990 and published by Kluwer Academic Publishers, Boston.

Remember, when reading these texts on communications, that (while the principles involved need to be understood) the distances required and the error rates obtained make transputer links much easier than telecomms.

A great deal of development is taking place in fibre connections, and probably the easiest way to keep in touch with the developments is by taking magazines, such as *Lightwave* or *Laser Focus World*, both from PennWell. More technical is *IEEE Lightwave Communication Systems*.

A good introduction to fast, low cost, optical fibre connections is given in HP's Application Bulletin 78, document 5954-8478 (3/88).

A number of standards are mentioned in this chapter, including SCSI and HPPI which are parallel interfaces, RS232, Ethernet, and Token Ring which are copper cable based LANs, and FDDI, FibreChannel and SONET which are optical fibre standards for LAN, computer interface, and long-distance telecomms respectively. After these standards are formally issued, they may be obtained from the standards authorities such as ANSI and IEEE. Obtaining drafts before the standards are published is not always easy, and may require contact with the working group responsible for the particular standard.

4.10 Manufacturers and products referred to

AT&T: 41 series of high performance line drivers, receivers, and transceivers;

Hewlett Packard: 820nm low cost 150Mbits/s fiber optic LED and receiver modules;

Honeywell: 820nm low cost 150Mbits/s fiber optic LED and receiver modules;

Madison Cable: 'SCSI' type cable with specified and low skew.