AVIONIC SYSTEMS STANDARDISATION COMMITTEE
NETWORK SYSTEMS SUBCOMMITTEE

Doc No : - ASSC/120/2/77 *

Issue 1 - June 1999

GUIDE TO HIGH SPEED INTERFACE STANDARDS*

0 EXECUTIVE SUMMARY

The purpose of this document is to give introductory guidance on network applications for future avionic systems by providing information on emerging standards and their current status. The document is intended as a guide for MoD Desk Officers and Industry on what choices of interface standard are available and inform them of development relating to avionics. The standards under consideration were selected by the working group and consist of the following:

Asynchronous Transfer Mode (ATM)
Scalable Coherent Interface (SCI)
Fibre Channel (FC)
Serial Express (P.2100, previously P1394.2)
Synchronous Digital Hierarchy (SDH)
Firewire (IEEE 1394-1995)
Ethernet and Gigabit Ethernet
Recs. X.34 & X.39:1996 - Data networks and open system communications, Public data networks – Interfaces

This general subject is advancing rapidly in the commercial field and across many applications areas. Therefore this document is intended to be an introduction to those current commercial standards that are possibly relevant in the area and to provide some information on the technologies they adopt eg what the standard is, and what should be consider in applying the standards to avionics.

The report provides an overview of the communications revolution that is taking place everywhere which is intended to illustrate the problem of transferring best commercial practices into candidates for the corresponding task in weapons platforms. In contrast the

---

* This report is published by the Avionic Systems Standardisation Committee (ASSC) to advance the role of standardisation in avionics. The use of this is entirely voluntary, and its applicability and suitability for any particular use, including any patent infringement arising therefrom, is the sole responsibility of the user. Copies of this paper are obtainable from the ASSC Agency as an Avionic Systems Standardisation Committee publication.

text further illustrates that, despite the best of intentions, there is no visibility of an imminent, clear cut Commercial off-the-shelf (COTS) solution that solves most problems.

This report subsequently takes a number of standards in turn and outlines the structure of each protocol, seeks to identify their strengths and weaknesses and proposes the types of system for which each standard is most applicable. Each standard tends to have been developed for a specific application area which may limit their use in other applications.

Table of contents **Page no.**

Table of contents (cont) **Page no.**

Figures and tables

1 INTRODUCTION

1.1 Need for standard interfaces

The principal military data communications standard in use throughout NATO is US MIL-STD-1553B otherwise known in the UK as Def Stan 00-18 (Part 2) - Serial, Time Division, Command/Response Multiplex Data Bus and in NATO as STANAG 3838 AVS. This standard was finalised before 1980. New applications need 'state of the art' performance of weapon systems and depend on computational devices capable of providing and requiring data that is orders of magnitude faster than US MIL-STD-1553B can provide. In recent years in the absence of any general solution, expedient solutions have been adopted to deliver design performance for particular functions. It is widely believed that a general strategy should be developed to gain all the well-known benefits of scale and commonality. This document is a step along that route in that it is designed to illustrate some of the problems and describe the features of some of the contending solutions.

In the late eighties there were three obvious contenders to be a long-lived military standard. These were adoption of FDDI, or the military initiatives in SAE i.e. High Speed Token Passing Ring Bus and Linear Token Passing Multiplex Data Bus. In less than a decade these have been substantially overtaken by several civil initiatives that were barely visible in 1990.

1.2 COTS principles and problems.

With the changing political scene and the enormous momentum of civil telecommunications the notion of transferring civil communications standards into military applications (COTS) is a widely understood principle. This principle is simple to conceive but can be difficult to deliver.

For many individual military communications problems there is a comparable civil problem with a corresponding solution. The well known time critical and high integrity problems of controlling a weapons platform are shared by many civil applications e.g. highly exothermic processes in petro-chemical plant and power generation, e.g. nuclear plant. These civil processes are currently being implemented by systems that are comparable in speed and integrity with US MIL-STD-1553B. A Guide to low and medium speed interface standards is included in another ASSC publication (ASSC/110/2/42). New technology is used to deliver new services. To a first approximation the future need is for high throughput for more or less permanent channels. This means that the contenders should be sought from comparable civil applications e.g. backbone telecommunications, computer backplanes and or video.

1.3 Speed

The advent of fibre optic technology has created a situation where absolute interconnect speed is unlikely to present a problem. To a first approximation all the protocols under consideration could deliver data throughput in excess of immediate requirements. It is also very probable that orders of magnitude of speed increase could be provided by advances in modulation/demodulation techniques in a far future refit. For COTS in the immediate future the problem is not speed but a mix of issues such as lifecycle costs, durability, integrity, vibration, repair, life expectancy.

1.4 Efficiency

The standards where transport inefficiencies are exhibited by the transfer protocols they specify. These are concerned with padding packets out to standard lengths or the ratio of overhead to payload. When making decisions to adopt a particular standard, academic efficiency considerations should be balanced with consideration of the whole problem. The whole problem can become very complex. A simple analogy is offered by an analogy to freight containerisation. If the cargo is sufficiently valuable and urgent it is quite meaningless to worry over empty space in the container when there is no economic alternative to transport the goods.

The other efficiency issue is the ratio of packet overhead to payload. All systems need management. In a climate where bandwidth is inexpensive, for the foreseeable future, it becomes economic to trade notional efficiency for other benefits notably, smaller granularity, technical simplicity and so on. Efficiency issues are easy to understand when presented in everyday terms but when the same issues arise in network protocols, they become more difficult to grasp.

1.5 Unified network

A unified avionics network could be implemented using a judicious adaptation of COTS protocols. Academically that is probably true. The advances in fibre optic technology and very high speed switching fabrics are being used to deliver novel services at a premium price or renew the old long haul backbone services that were once coaxial cable. For most people the telephone cable into the home might be 20-30 years old (or more) but it is not economic to replace the twisted pair for fibre optics. In other words there is usually no economic advantage in replacing an obsolete technology where it serves the current purpose well enough. New technology is used to deliver new features. This is true in the civil arena where cost is a more visible driver but the military circumstances will not be very different.

All of this is to suggest that the shift to a unified network for avionics is likely to be a gradual process over a long term. Amortised solutions using US MIL-STD-1553B will remain an economically viable solution for a long time by simple existence of tried and tested functions particularly if its transfer rate can be increased without having to replace installed bus cabling or components. A step function to a complete new network architecture has risks and costs that are probably unaffordable but also unjustified.

1.6 Fitness for purpose

To a first approximation all the following protocols are likely to be similar solutions to the basic technical problem. Where the problem is specified in terms of throughput, latency and other networking parameters then it should be possible to prepare a marking scheme and so score these parameters to create an order of preference. Unfortunately the mainstream civil applications do not share certain mechanical stresses that are commonplace on a weapons platform. These are shock, temperature, vibration, EMP, flammable atmospheres, gunfire, battle damage, maintenance, bulkheads and so on. It is arguable that the principal uncertainties lie more with the physical properties of fibre optic interconnects than with the protocol consequently the problem is generic. In other words each COTS prospect has to be evaluated for military requirements especially those that do not have an equivalent in the original civil application. This evaluation is considerably more than a desk exercise.

It is apparent in civil telecommunications that the same physics is being used to construct different applications. This means that Giga-bit Ethernet is specialised to address high speed LAN-LAN Ethernets, ATM is for long haul international telecommunications. These are not head-on competing ideas but market specialisation of the same principles. The various civil applications have special needs that give rise to specialised solutions. Military applications also have special needs that are not really satisfied by any civil initiative.

1.7 Cost of ownership (lifecycle costs)

It is arguable that any of the identified network protocols could meet the throughput requirements for the foreseeable future. This indicates that best choice is dictated by other factors. Where technical merit is not a performance issue then cost is almost certainly the deciding factor. Real comparative costs are notoriously difficult to estimate. The life of an airframe is perhaps 30 years whereas the useful lifetime of some civil protocols might be 5 years, even less. Due consideration should be given to knowledge base, maintenance tools, complementary protocols, installed base, stability, level of support, etc.

Every prospective network protocol is not an island but exists in a commercial context of complementary functions. The need for COTS does not stop with just the choice of network protocol. It is the availability of all the complementary functions that most affect the long-term total cost. For example, the driver for more bandwidth in civil telecommunications is video in one form or another. The result is close interdependence of specific application (e.g. a video protocol) with a specific delivery service (network). The high bandwidth network is specified to bear video, the video protocol in turn is optimised to interface well to a particular bearer. In the fullness of time all the components in a chain become available off the shelf. Conversely, it is likely that transporting applications on some other network might be anywhere between unfeasible and expensive. The cost reduction from applying COTS principles has to be estimated from the total cost of the aircraft not just decided on the costs of the network protocol.

1.8 Standardisation

Traditionally standards have been created within institutions such as ISO, IEC, ITU, and IEEE. Within the IT and telecommunication industry the ‘time to market’ and ‘market lifetime’ has resulted in huge pressures. As a result a new route to standardisation has emerged in the last few years and this trend is accelerating. This is ‘market led’ standardisation visible in ATM forum, Gigabit Ethernet Alliance, Fibre Channel Association and so on. These fora exist briefly, to develop industry consensus and to share out the work and then be dissolved. The cost of membership is set high to limit attendance to serious players. The standardisation is run as ‘a project’ using normal project management techniques. The net result is optimum convergence on the consensus solution to a particular initiative. ‘Fast track’ submissions are scheduled to deliver ‘finished standards’ into the traditional standards bodies e.g. ISO or IEEE. For this text the point being made is Giga-bit Ethernet is a project driven by industry however, the finished work will be endorsed and published by IEEE. The result is that a proportion of industry practice is embodied in ‘the standard’ but crucial ‘implementation agreements’ can exist outside. A further result of this trend to market led standardisation is that full access to the drafts and processing has a price.

1.9 Overview summary

The purpose of the above discussion is to illustrate that a strategy for identifying a common policy for a future network protocol has major issues that require careful investigation. No single project can afford investigation beyond the needs of that project. In other words to substitute a calculated policy for the current expediencies will require dedicated work. Even so, it is further stressed that evaluation of the prospects depends on a valid set of operational requirements.

In summary, all the network protocols under consideration have been established in the nineties and the rate of progress in network protocols and novel applications e.g. digital video is probably accelerating. In this climate the notion of a long-term strategy endorsing a single high speed network protocol is probably abortive. The safest course would be to select a short-list, possibly the list that follows, and set about the necessary work to identify the general fitness for use in avionics but leave a final choice to be project specific.

1.10 Objective

It is now widely recognised in the field of military avionics that as processing power increases and sub-systems evolve, the absolute performance of the platform is dependent on an increasing need to transmit yet more data which will be a key element in future avionics architectures. It will also be necessary to ensure that latency values for the data transmission are within the required deadline, and that failure to meet such a deadline is considered a system fault which may result in the failure of the mission. Different protocols are been defined for each principal civil application, each of which have their own strengths and weaknesses, making some better suited than others to certain avionic system configurations.

The aim of this section is to highlight the different factors associated with each of the standards and to identify applications for which they are most suitable. The report can thus be used to aid the process of choosing the appropriate standard for a particular application or determining the attributes of standards being offered.

Not all systems will benefit from higher data transfer rate because the architecture used may reduce the total amount of data that is transmitted over the system or because no increase in the total data is perceived to be required during the lifetime of the system. Factors which must therefore be taken into consideration when determining which standard is best suited to the intended application include the performance of the standard; implementation issues such as cost and complexity; and the expectation that established standards will require less development time than one that operates an entirely new protocol. As processing and memory requirements become even greater, so too do the requirements associated with the data transfer rate to ensure throughput remains at the level required by the system.

The layout of this report is such that a general overview of each of the data buses is given, followed by a section giving details of the strengths and weaknesses. Also included is a section describing the main applications of each data bus. Appendix 1 containing a table which summarises the main features and differences of the protocols is given at the end of the report. The description given for each standard is that of its conventional usage. It is noted that all standards can be modified to meet particular system requirements (e.g. through the addition of protocols which use the data bus transfers), but doing so may result in transfer inefficiencies, incompatibility with other systems using the same data bus and problems not previously encountered which may make commissioning as difficult as for a new data bus.

2 DESCRIPTION OF DATA BUSES

2.1 Asynchronous Transfer Mode (ATM)

2.1.1 Background

Asynchronous Transfer Mode (ATM) was initially conceived, in the late 1980s, as an integral part of the ITU-T's[1] (formerly CCITT) B-ISDN[2]. B-ISDN is to be the world-wide switching fabric for telephony and data and will eventually carry all short duration traffic. This includes telephony LAN-LAN data, video and is also a packet switch network supplanting X.25. B-ISDN is intended to extend ISDN beyond its current, relatively low data rates[3] to the much higher data rates envisioned for the next generation of fibre-based telephony[4].

---

[1] International Telegraph Union (ITU) Telecommunications Standards

[2] Broadband Integrated Services Digital Network

[3] ISDN offers interfaces in the low kbps to low mbps range. ISDN is defined in such a manner that it operates effectively over the plesiochronous data hierarchy (PDH) that defines the meaning of DS0, DS1, etc.

[4] B-ISDN is defined to be commensurate with the Synchronous Data Hierarchy (SDH), which defines synchronous data channels in the 155 Mbit/s to 2.5 Gbit/s range.

The reader should note that ATM is a packeting and switching technology which is rate independent (i.e. scaleable). ATM requires a companion physical bearer service. Within ITU and B-ISDN that service is SDH Synchronous Digital Hierarchy. The actual transmission networks now being built world-wide are SDH (The UK Energis backbone network is entirely SDH) but the synchronous data streams for such as telephony and FAX, etc. will be eventually assembled by ATM from lower speed bursty subscriber data streams.

A useful analogy for comprehending ATM (and its relationship with SDH) is to liken the resulting data network with freight containers. ATM manipulates the application data into a succession of packets (containers). SDH is the backbone transport network. At the point of use, ATM packets are analogous to containers submitted to or removed from the freight network. Every container carries a reference to a final address and every user is given a delivery schedule. At each hub (switch) along the whole journey the switch routes containers with a common address and dispatches them accordingly. In other words the transport system (SDH) has no care for content only routing and attributes (e.g. priority). The pure idea is a global system of ATM propagated on SDH. In reality, like with the freight analogy, the containerisation (ATM) might exist end to end but SDH is supplemented by fast Ethernet, FDDI, even twisted pair telephone using Asynchronous Digital Subscriber Loop (ADSL). Additionally SDH is a hierarchy of bearers. Each faster bearer is scoped to carry exactly four tributaries.

In civil systems. ATM is no more a complete solution to all problems than it would be for military systems. Any high speed permanent connection, such as between video studio and transmitter, will probably be transported directly by a leased SDH circuit. Correspondingly, in any avionics application that requires permanent high throughput connection could use SDH directly. The analogy here is that it is not optimum to use freight containers to deliver coal to a power station.

ATM started out as a wide-area switching technology, but many soon realised that it could be employed in the local area as well. Indeed, one of the most attractive potential advantages of ATM is that it represents an opportunity to integrate the wide-area, metropolitan-area, and local-area domains with a single, "seamless" network - a technology, moreover, that promises to integrate key classes of digital services as well.

Asynchronous Transfer Mode (ATM) is a flexible technique for which a network can be designed that will be capable of expanding to meet the requirements of future services or applications. This is the major advantage of ATM over conventional packet switched networks and it is being viewed as the means to make the telecommunications Broadband Integrated Services Digital Network (B-ISDN) possible. B-ISDN is a concept that began to be evolved in the late 1970s to provide a global strategy to meet future communication requirements. It was foreseen that the sheer volume of information transfer would grow rapidly and that an integrated approach would be needed that provided a wide range of services from voice through data to video. These services could only be provided on a global scale by the use of digital technology.

2.1.2 Description

ATM is a fast packet switching/multiplexing technique, (it uses short, fixed length packets or cells as the common means to support all services) in which packets are switched across an SDH circuit(s), from ATM switch to ATM switch, until they arrive at their destination. ATM is service independent and bit-rate independent and when proposed for use in telecommunications applications uses a statistical allocation and sharing of available resources.

Common to all technologies that bear the name of or associate themselves with ATM is the notion of *cell-relay*. A cell is just a fixed-size, fixed-format packet. The ATM cell is 53 bytes long, of which the first 5 bytes are *header* and the remainder (48 bytes) *payload*. The header is mainly used to determine the routing of the cell through a network.

A variety of types of service can be transported by ATM and to accommodate these an adaptation function is provided to fit the information into ATM cells and to provide service-specific functions. ATM will transport both fixed and variable bit rate services and also offers two priorities: high for cells with negotiated throughput, and low priority, with the possibility of the cell being discarded, if the resources are not available.

2.1.3 Advantages and disadvantages

The principal advantage of ATM is that it is aimed at using network resources in the most efficient manner possible (that is to match fixed network resources to variable subscriber demand). The network's role is to deliver packets of information and it does this by using the five byte header, the network is not concerned with the payload itself. The fixed cell size and the transfer of payload error checking to the endpoints of the connection allow switching and buffering to be easily performed in hardware and provides minimisation of padding of packets to an integral number of cell payloads.

Among the advantages of such a rigid packet structure are

- Ease and lower cost of implementation of cell processing in a VLSI chip
- Higher per-packet processing speed
- Lower per-packet queuing delay
- Easier buffer allocation

The key disadvantages are bandwidth inefficiency (there is an unavoidable 9.4% bit-overhead in ATM) and processing overhead for many types of traffic (e.g., bulk data transfers) that are more suited to larger, variable-size frames. The relatively small size of the ATM cells contributes to lower and less variable network latency and easier allocation of bandwidth, but exacerbates the disadvantages mentioned above for some traffic types.

ATM has the potential to offer error- free routing of information by virtue of the single byte of error control information which can provide detection and correction to the header. (It is believed that this provides protection against a single bit error.) In normal data use the real error correction abilities are provided by higher layers of protocol. That situation will not be different in principle for military systems. It should be noted that ATM is unidirectional so two channels will usually have be needed.

ATM can provide an effective transport mechanism for constant bit rate data, e.g. video, while at the same time providing variable bit rate services, e.g. message passing. The decisive advantage of ATM/SDH over other network protocols is the rapidly accelerating pace of deployment and ubiquitous attachment of novel applications, e.g. video in the form of MPEG-2 developments.

At the present time, commercially available switches only operate up to 622.08 Mbit/s although 1.2 Gbit/s switches are soon to be available. There are, of course, no switches that have been specifically designed for avionic use. The greatest disadvantage of ATM is the so called admission problem. For every new subscriber a civil network has to calculate the future ability to fulfil 'contracts' made with all existing connections and the new subscriber. Fortunately this problem need not arise in a closed environment of an airframe where the absolute worst case can be calculated in advance and guaranteed capacity deployed.

2.1.5 Cell loss; cell discard

As with any packet switched network, overloading the switches in an ATM network will result in cell loss due to congestion. Whilst this is considered acceptable in the commercial applications, it will be necessary for congestion to be avoided in military avionic systems. Since in this case, it is possible to have a full description of the traffic, there is no need to rely on statistical models of the sources of data. Hence, it is possible to predict whether congestion can or cannot happen in a given implementation. This will allow the calculation of worst case delays, necessary to the implementation of real-time applications with hard deadlines.

2.1.6 Risk

ATM is still a relatively new technology and as such it presents some risk. The physical bearer, SDH is substantially complete and being widely deployed (The UK Energis backbone network is SDH). However, note that whilst ATM is bound to SDH in the B-ISDN, in practice ATM could be constrained to use other transmission networks. There is still standardisation work being undertaken by the ITU-T and the IEEE has started an ATM forum. This constitutes some risk if protocol and component development do not develop to include avionic specific features.

2.1.8 Standards

A number of organisations are involved in standardising ATM/SDH:

- Internationally, the ITU-T (formerly CCITT)
- In the US., the American National Standards Institute (ANSI)

- Among ATM vendors, the ATM Forum
- For inter-operation with the Internet community, the Internet Engineering Task Force (IETF)

2.1.9 www pages

There is very extensive information available on ATM, an easy way is browsing www.
An excellent starting place with many links is
http://mailme.hill.com/library/atm_overview.html

2.1.10 Conclusions

Asynchronous Transfer Mode (ATM) offers a flexible networking solution for all types of services from low to high data rates, fixed or variable bit rates. The use of ATM for avionic networks depends mostly on switching element availability and switching element size. Currently ATM switches have a bit rate per input of 155 Mbit/s but this is likely to increase with LSI and VLSI in the future. European collaborative programmes are studying optical switching for ATM which would increase the throughput bit rate tremendously.

The possible problem of cell discard in ATM switches may be overcome in avionics by dimensioning the switches/interconnect adequately to handle the worst case % loading. It would be possible to dimension the switches/multiplexers and interconnect for worst case % load by virtue of the fact that an avionic system has a fixed connectivity and known worst case data rates for each and all paths. Cell loss due to header errors can be kept to a minimum by use of an adaptive detection/correction method, but this may need to become a full forward error correction scheme to ensure no cell loss. Commercial ATM networks can offer Quality of Service guarantees including bounds on cell propagation delay. Not all commercial techniques for providing such guarantees are adequate for real-time data transmission in avionics systems.

Information about the size, capacity and delays associated with ATM switches, is necessary to fully determine the extent to which ATM is a suitable transfer mode for avionic type networks. In civil application ATM is designed for huge numbers of short overlapping connections with dynamic addressing. When connection time becomes extended (many days) and data rate is high, e.g. High resolution MPEG-2 video, then ATM serves no purpose since it will be better to access an SDH point to point route. The same might be true for an avionics platform, i.e. ATM is a good vehicle to aggregate and route all the low rate services but the high rate services are applied directly to the physical bearer and switched on a ‘hard wired’ route.

An ATM network offers maximum flexibility for an acceptable complexity. The flexibility is visible to the users as an unlimited arrangement of the generated bit rate. The bit rate may vary from a few bits per second up to 150 or 600 Mbit/s with any value in between, and may be a fixed or variable bit rate. Users may also operate at a clock speed which is independent of the network clock.

2.2 Synchronous Digital Hierarchy (SDH)

2.2.1 Background

A concerted standards effort has been involved in the development of synchronous transmission . The opportunity of defining this new standard has been used to address a number of other problems. Among these have been a network management capability within the hierarchy, the need to define standard interfaces between equipment, and European transmission hierarchies.

This standards work culminated in CCTTT Recommendations G.707, G.708, and G.709 covering the Synchronous Digital Hierarchy (SDH). These were published in the CCTTT Blue Book in 1989. In North America ANSI published its Synchronous Optical Network (SONET) standards, which can now be thought of as a subset of the world-wide SDH standards.

In addition to the three main CCTTT recommendations, a number of working groups were set up to draft further recommendations covering other aspects of the SDH, such as the requirements for standard optical interfaces and other standard functions.

The CCTTT recommendations define a number of basic transmission rates within the SDH. The first of these is 155 Mbit\s, normally referred to as STM - 1 (where STM stands for 'synchronous Transport Module'). Higher transmission rates of STM - 4 and STM - 16 (622 Mdit\s and 2.4 Gbit\s respectively) are also defined, with further levels proposed for study.

The recommendations also define a multiplexing structure whereby an STM – 1 signal can carry a number of lower rate signals as payload, thus allowing existing PDH signals to be carried over a synchronous network.

SDH defines a number of "Containers", each corresponding to an existing plesiochronous rate. Information from a plesiochronous signal is mapped into the relevant container. The way in which this is done is similar to the bit stuffing procedure carried out in a conventional PDH multiplexer. Each container then has some control information known as the "path overhead" added to it . The path overhead bytes allow the network operator to achieve end-path monitoring of things such as error rates. Together the container and the path overhead form a "Virtual Container".

In a synchronous network, all equipment is synchonized to an overall network clock. It is important to note, however, that the delay associated with a transmission link may vary slightly with time. As a result, the location of virtual containers within an STM - 1 frame may not be fixed. These variations are accommodated by associating a pointer with each VC. The pointer indicates the position of the beginning of the VC in relation to the STM - 1 frame. It can be incremented or decremented as necessary to accommodate of the position of the VC.

G.709 defines different combinations of virtual containers which can be used to fill up the payload area of an STM - 1 frame. The process of loading containers, and attaching overhead is repeated at several levels in the SDH, resulting in the "nesting" of smaller VCs within larger ones. This process is repeated until the largest size of VC (a VC - 4 in Europe) is filled, and this is then loaded into the payload of the STM - 1 frame. (This subject will be discussed in more detail in Chapter 4). When the payload area of the STM - 1 frame is full, some more control information bytes are added to the frame to form the "Section Overhead". The section overhead bytes are so-called because they remain with the payload for the fiber section between two synchronous multiplexers. Their purpose is to provide communication channels for functions such as OA&M; facilities, alignment and a number of other functions.

When a higher transmission rate than 155 Mbit\s of STM - 1 is required in synchronous network, it is achieved by using a relatively straightforward byte - interleaved multiplexing scheme. In this way, rates of 622 Mbit/s (STM - 4) and 2.4 Gbit/s (STM - 16) can be achieved.

2.3 Scalable Coherent Interface (SCI)

2.3.1 Background

SCI is an IEEE approved standard (IEEE Standard 1596-1992) intended to be the next generation high speed computer backplane for interconnections in multiprocessor machines. SCI was designed to use point-to-point links in order to avoid the physics problems in using a backplane transmission line at very high data rates, e.g., distributed capacitances. A distributed shared memory structure is used by SCI to maintain cache coherence. SCI also supports message passing in addition to shared memory transfers.

Backplane computer bus approaches have several characteristics which make them unsuitable for very high speed processes. Only one data transmission can be handled at a time in a backplane bus, thus the bus can become a bottleneck for a multiprocessor system. Since a backplane bus is actually a transmission line, the limiting effects of such factors as capacitance and propagation delay become more pronounced as higher frequency signals are passed on the bus. Similarly, the effects of reflections due to improper impedance matching at connectors as well as different loading characteristics (depending on how many modules are inserted in the backplane) can severely limit bus performance.

The stated purpose of SCI is "to define an interface standard for very high performance multiprocessor systems that supports a coherent shared-memory model scaleable to systems with up to 64 K nodes. This Scalable Coherent Interface (SCI) standard is to facilitate assembly of processor, memory, I/O, and bus adapter cards from multiple vendors into massively parallel systems with throughputs ranging up to more than 10^{12} operations per second." [5], [6]

---

[5] IEEE SCI Std (P-1596)

[6] 64 K implies 64,000

SCI uses point-to-point signalling to simulate a bus without actually using one. This results in higher speeds due to greatly simplifying electrical transmission problems. In order to keep track of how the data is being used (e.g. has it been successfully received, whose turn it is to transmit, etc.) protocols have to be used which are different from those used with buses.

SCI is emerging as a commercial interface, with several chip manufacturers developing or manufacturing products. It has also received significant attention from the military, particularly in the US, where SCI has been chosen as the baseline for the US Navy Next Generation Computer Resources (NGCR) High Speed Data Transfer Network (HSDTN). SCI has also been selected as the baseline for the Society of Automotive Engineers (SAE) Sensor Video Interconnect Network. This military interest is particularly concerned with real-time systems as opposed to the fairness schemes used in many commercial systems. This has prompted the development of an SCI/Real-Time Group P1596.6, who are focusing on adding mechanisms to allow SCI to meet real-time requirements.

2.3.2 Description

SCI operation is based on ringlets, and it provides for great flexibility in implementation topology. High performance systems are envisaged as being constructed from switches (Figure 10), with the processors being attached to the switch by two-node ringlets. Low cost systems are seen as using larger ringlets, with bridges to other ringlets for more connectivity.

Basic protocol operation is by means of request and response sub-actions. These sub-actions are treated separately in the interface to prevent deadlocks occurring. Request and response sub-actions consist of a header consisting a number of 16 bit words for source Id, target Id, command and control. Addressing of 48-bits (in addition to the 16-bit node address) can be provided to a data payload of 0, 16, 64 or 256 bytes. Finally, a 16-bit CRC based on the "CCITT polynomial" is appended to the packet. Request and response sub-actions are followed by an echo, in which the receiving station acknowledges the request or response.

16-bit idle packets are used to fill the spaces between transactions, these provide synchronisation between transmitters and receivers and help maintain the ringlet access mechanism. Init. packets are used during ringlet initialization.

Being based upon a point-to-point topology. SCI broadcasting is more difficult than in a shared media bus system. Broadcast is provided for by the broadcast move transaction, which has an echo, but no response sub-action is expected. The broadcast receive capability of a node can be disabled, providing the possibility of a multicast scheme to be used. However, this is outside the scope of the base SCI standard.

SCI also contains a cache coherency mechanism for support of multiprocessing. This mechanism is useful in a non-shared medium system to enable all local copies of data to be kept synchronised (or coherent). This is achieved using a doubly linked list of locations of cached data which is used to send updated messages to those locations with cached data.

2.3.3 Performance

SCI has several media options for its point-to-point link-s including 1 Gbyte/s parallel electrical, 1 Gbit/s serial electrical, 1 Gbit/s serial single-mode fibre optic, and Low Voltage differential options under development by the P1596.3 working group.

The single serial link runs at 1.2 Gbaud, using the 16B/20B Hewlett Packard CIMT (Conditional Invert, Master Transition) block codes. If 256 byte packets are used, a 14 byte header, and a 2 byte CRC, then the overhead is 6.25%. This is a minimum figure since in practice the insertion of idle symbols (at least one per packet), and the sending of 8 byte echoes will reduce the bandwidth for data.

With the variety of media and topology options, it is hard to assess the performance of a general SCI system. The maximum bandwidth per node will be 8 Gbit/s, although a parallel switch operating at this speed may prove very difficult to implement. It is therefore likely that parallel implementations will be using ringlets, with the 1 Gbyte/s bandwidth shared amongst the nodes. A 1 Gbit/s switched system may prove feasible in the near-term. SCI has 16-bit physical addressing capability, so that it is likely that switched system throughput will be limited only by the size of the switches available.

The SCI ringlet protocols allow multiple concurrent messages on the ring, increasing the bandwidth available per node. Bandwidth allocation protocols are used to inhibit the transmissions of some nodes to ensure transmission opportunities for others. Two levels of priority are provided, fair (i.e. low priority) which gives equal access to the ringlet with no node preferred over any, other, and unfair (i.e. high priority). Access to these priorities is provided independently by means of low-go and high-go bits in the idle symbols. Messages are placed on the ringlet by stripping of the idle symbols, and the protocol ensures access for all nodes by maintaining the appropriate balance of idle low-go and idle high-go symbols. It is this bandwidth allocation protocol which is receiving most attention from the SCI/Real-Time group, who aim to sacrifice the fairness mechanism to obtain real-time performance. This proposed standard (IEEE P.1596.6) is developing improved support for real-time and multimedia applications of SCI, reduced redundancy for priority traffic. Fault tolerant mechanisms for SCI are also being studied.

Queue allocation protocols are also used to help improve performance. Mechanisms are included to reserve queue space in a node which is waiting for a response. These protocols should also prevent deadlock from occurring. Separate queues for requests and responses are specified in SCI nodes and switches to help with this problem.

The SCI link is estimated (generically) to be able to transfer 378/471 Mbit/sec (global/local, using one 64 byte and one 16 byte frame) of messages or 936 Mbit/sec for data streaming transfers. Multiple, independent links in the system, many operating in parallel at any one time, allow enormous potential throughput. A module input/output is limited to 1.8 Gbit/sec (half input, half output) through a single interface.

2.3.4 Implementation

It is unlikely that the basic SCI standard would be implemented in an avionic network. The SCI/Real-Time modifications should provide the necessary features to enable its use in avionic systems. It is likely that a variety of other modifications (such as to the media), will be needed to make the components suitable for the environment. SCI switches may require additional hardware when compared to an ATM switch, since each input and output must implement an SCI node. SCI switches are also anticipated to contain separate queues for request and responses, perhaps doubling the number of buffers required.

The topology options of ringlets and switches could vary between two extremes. A fully interconnected system could be built, where each module has an interface to a switch, and the ringlets contain just two nodes. At the other extreme, ringlets could be used within the racks, which could have just one interface to a small dimension switch. Within the latter system, ringlets might be built out of 1 Gbit/s serial links which are more suitable for optical waveguide implementation, or could use the 16-bit parallel option for increased bandwidth. The choice of topology will be influenced by the communication requirements of the modules, the need for reconfiguration outside the racks, capability for growth, etc.

2.3.5 Risk

SCI shares with other comparable interconnection systems, the risk associated with high speed fibre optic links, laser stability, etc. SCI node commercial components are now beginning to emerge, and there does appear to be considerable development activity occurring (although at a lower level than ATM). The SCI/Real-Time extension, which is more likely to be implemented in an avionic system is still very much in its infancy, and as such must represent a risk. The lack of consideration of SCI switches must also represent a risk.

2.3.6 Conclusions

SCI offers a high performance interconnection system for multiprocessing. It is this performance, coupled with its inherent support of the mechanisms required by processors, which will enable it to become an important standard in the commercial computer arena.

It is anticipated that P.1596.6 (SCI/RT) will enable a modified SCI system to meet real-time requirements. It may be that such a standard will prove suitable for high performance avionics backplane type applications. It remains to be seen however, whether SCI is suited to carrying streaming data such as sensor video traffic. At the least, the features and mechanisms of SCI will be useful in the assessment and development of other spatial switching systems to be used for processor interconnection.

2.3.7 Standards

There are currently six SCI official IEEE follow-on efforts underway. For example, SCI/Real-Time is an SCI follow-on group which is investigating the use of SCI protocols for real-time applications which require guaranteed latencies. SCI/RT also provides some increased fault tolerance and error handling capabilities. There are also several study groups which are investigating various extensions and enhancements to SCI. These are :-

P1596.1:	Guide for Switches/Bridges.
P1596.2:	Kiloprocessor extension. (developing efficient cache coherence mechanisms for use in very large SCI systems).
P1596.3 :	Low-Voltage Differential Interface for the Scalable Coherent Interface (also defines 4-, 8-, 32-, 64-, and 128-bit wideSCI links).
P1596.4:	RamLink, defines a high speed memory and I/O interface based on point-to-point links.
P1596.5:	Shared data formats (shows how to describe shared data).
P1596.6:	SCI for Real Time Applications (SCI/RT)
P1596.7:	High Speed Memory interface (formerly called ‘synclink’ now promoted by a consortium as ‘SLRAM’)
P1596.8:	Working group to detail cables and connectors for SCI applications
P1596.9:	Application Program Interface for SCI

2.4 P2100 SerialExpress: (formerly P.1394.2)

2.4.1 Background

A high performance low-latency scalable interconnect, with no particular limits on bandwidth or distance, suitable for both consumer and industrial applications, such as digital video and audio streams, computer peripheral I/O, computer cluster interconnect, NUMA and ccNUMA multiprocessing.

SerialExpress is designed to make switch hubs and bridges to other standards or to itself as simple as possible, and to support multiple concurrent transfers so that system bandwidth can be expanded, essentially without limits, by adding cables, bridges, and switches.

SerialExpress uses small packets for low latency and for keeping packet buffers inexpensive. "Serial" cables transport these packets one bit at a time, but the same protocols work with wider links as well, which is especially economical for short links (up to about 10m) and for applications that require very high bandwidth (say 500 MBytes/s to multiple GBytes/s per cable).

2.4.2 Overview

SerialExpress is the merging of principles from SCI (IEEE 1596) and the original IEEE 1394 called ‘Serial bus’ (also called ‘Firewire’ by Apple). A brief description can be found from http://www.scizzl.com/P1394.2/index.htm.

SerialExpress is undergoing intense development fuelled by SUN, Intel, Apple and many others. The market place is work-station, PC backplane and peripheral video connections.

2.5 IEEE 1394 Serial Bus standard (Firewire)

The IEEE 1394 (also IEC 61883) data transmission protocol provides high speed, low cost, and a user-friendly interface. It has 100, 200, and 400 Mbits/s rates

The IEEE 1394 multimedia connection enables simple, low-cost, high-bandwidth isochronous (real-time) data interfacing between computers, peripherals, and consumer electronics products such as camcorders, VCRs, printers, PCs, TVs, and digital cameras.

The 1394 digital link standard was conceived in 1986 by technologists at Apple Computer, who chose the trademark 'FireWire', in reference to its speeds of operation. The first specification for this link was completed in 1987. It was adopted in 1995 as the IEEE 1394 standard. A number of IEEE 1394 products are now available.

The 1394 Trade Association (http://www.1394ta.org) has separated the extensions into individual study groups.

IEEE 1394.1 is the supplement to 1394
IEEE 1394.2 is the gigabit extension
IEEE 1394a is the Standard for high performance serial bus (a new standard project)
IEEE 1394b is the standard for the Serial Bus high speed supplement (a new standard project)

This information does not entirely match yet more Intel data (http://developer.intel.com/solutions/tech/1394.htm) that purports to be reporting the same endeavour. It is apparent that this marketplace is characterised by marketing trade names not numbered standards.

2.5.1 Marketplace

This initiative now has large amounts of commercial noise and hyperbole. Products are anticipated in 1999 however, the nature of the core market suggests a short product life.

2.6 Fibre Channel (FC)

2.6.1 Background

Fibre Channel refers to a set of standards under development by the ANSI Fibre Channel committee, X3T9.3. Fibre Channel specifies a high-speed serial data channel that can connect nodes point-to-point or through a switch or switch network ("Fabric" in Fibre Channel jargon).

Fibre Channel began in 1987 as a proposed physical infrastructure for the Intelligent Peripheral Interface (IPI), a high performance disk/tape interface used in many high-end systems. Later, the scope widened, so that Fibre Channel was seen as a standard low-level transfer mechanism for carrying protocols such as HIPPI, SCSI, IPI, and other bulk-data transfer services. At the time, most bulk channels were very limited because of their large footprint, the short distances over which they could work, the limited connectivity they supported, and the speed at which they operated[7]. Fibre Channel was to overcome these limitations in a standard way. Work on the first standard began in earnest in late 1989. Among the requirements and goals for Fibre Channel were the following:

- Small footprint (implies serial channel)
- 2 to 10 kilometre operating distance
- Up to 800 Mb/s in payload
- Efficient operation over long distances
- Greater connectivity than existing multidrop channels (e.g., SCSI)
- Efficient multiplexing of multiple streams into a single port
- Support multiple existing interface command sets without modifications
- Support multiple cost and performance levels

As work progressed, the committee members (and others) came to realise that they were defining a very general transfer mechanism that could be used as a networking technology as well as a data channel. Thus, despite its roots as a peripheral channel, Fibre Channel has come to be perceived as a technology that would support the construction of high-performance local area networks.

2.6.2 Description

The Fibre Channel (FC) standards define topologies, fabric services and interface details with which to make up a system. It is intended to provide a high performance serial link supporting its own and other data network protocols (a 'Universal Pipe'). The topologies identified are point-to-point, fabric and arbitrated loop. There are three Port Types defined :

- Fabric (F_Port); for point-to-point and fabric topologies
- Network (N_Port); for fabric topologies
- Loop (L_Port); for nodes and fabric in loop topologies (NL Port and FL Port)

Both point-to-point and fabric topologies are predicated on a bi-directional link connecting an N Port with another N Port of an F Port. Three classes of connection for data transfers are defined:

• Dedicated (Class 1)	have sole use of the resources required to provide the connection with end-to-end flow control and guaranteed delivery order.

---

[7] Since that time, serial HIPPI and HIPPI switches have been defined (but not standardised). These two advances in HIPPI technology also overcome some of the limitations listed in this sentence.

•	Multiplex (Class 2)	have frame delivery notification, but no guaranteed delivery order, with buffer-to-buffer (link) and end-to-end flow control.
•	Datagrams (Class 3)	have no delivery notification or order guarantee with just buffer-to-buffer (link-) flow control.

There is also an option within Class 1, called Intermix, which allows Class 2 and Class 3 frames during an established Class 1 connection.

The Fabric is transparent to the connected devices (nodes), but it does not passively pass data (i.e. it is not just a switch). An FC Fabric is a 'black box' which interconnects multiple nodes with no Fabric topology or implementation identified, but with interfaces performing definite functions (the F ports). The fabric may restrict communication models and classes of service supported. For this assessment, a single network based on switches is assumed for all transfers with the suggestion that this be hierarchically implemented to allow some limitation of data scope (i.e. data does not get transferred further afield than necessary).

FC uses serial, fibre optic links at various defined speeds: 132.8125, 265.625, 531.25, and 1062.5 Mbaud. It uses an 8B/10B encoding scheme (although identified through a 5B/6B + 3B/4B designation) which defines 256 Data (D) Codes and 12 Control (K) Codes (11 Reserved). It also defines Ordered Sets of four-byte words which are used for Frame Delimiters (8 Start-Of-Frame and 12 End-Of-Frame). Primitive Signals (e.g. Idle and Receiver Ready) and Primitive Sequences (e.g. Port Indications).

Before data transfers are carried out, the N Port must Login to the local F Port and to the N Port(s) of interest to determine the relevant capabilities affecting transfer parameters. The frame addressing is defined in the standard as physically based, although sharing addresses across multiple N Ports (alias address) and a broadcast address are also identified. The frame data field is limited to 2112 Bytes by clock- accuracy, although by adding extra Idles in the data this may be extended.

2.6.3 Attributes

Switch Fabrics allow multiple paths to be active simultaneously, albeit with the additional switch hardware being required. Rings are cheaper, because no additional hardware is required, but can only support limited simultaneity. For the discussion, a switch based Fabric is assumed.

Since the transmitter output power exceeds Class 1 laser limits for some of the FC standard options, a technique to detect an open fibre link through the use of a laser control system is included. Since FC is a bi-directional link, the open fibre link detection ensures optical power is received, whilst transmitting little power, before providing continuous output power to the level defined for data transfer. If at any time the received input power is lost, the output power is immediately reduced again.

FC is optimised for large block transfers and includes segmentation and reassembly actions in order to transfer them within the restricted frame data field. These transfers are described in terms of a hierarchy of Exchanges, Sequences and Frames. Error detection and recovery are based on the links and Sequences. Acknowledgement of frames is performed end-to-end on either a Frame or a Sequence for Class 1 and Class 2 transfers, whilst link acknowledgement (F Port to N Port) is provided for Class 2 and Class 3 transfers. These are also used for flow control. An indication is provided from each switch in the Fabric and the destination node that the frame was/was not correctly routed. The likely reason for this apparent divergence is that a Fabric comprising circuit switches would normally be expected to mainly support Class 1 connections. It is also true that ring based Fabrics would not use such node by node acknowledgements.

Within the switch, frames are received from the fibre-optic receivers (FOR) and fed, still 8B/10B encoded to the Routing Protocol Decoding. Here the header is decoded, whilst also being saved still encoded in the FIFO. The protocol logic makes a connection request to the other Routing Protocol Decoding blocks and, if the required path is free, sets up the route through the N x N parallel crossbar to the required output and readies an acknowledgement to be returned to the sender when the relevant port is free. The whole frame may then be transferred from the FIFO using the transmitter clock. The idle at the end of the frame will then allow the FIFO to be reset.

If the required path is not free, the incoming frame is lost in the FIFO and a negative acknowledgement is readied to be returned to the sender when the relevant port is free.

2.6.4 Advantages and disadvantages

FC allows for a system of unlimited size by virtue of the Fabric not being defined. However, as indicated above, it is sensible to define a hierarchical structure to the Fabric to limit the scope of the data to just that necessary. It is suggested that a switch be provided within each rack with the switches connected together for inter-rack communications with as many links as is deemed necessary to prevent bottlenecks. No consideration has been given to redundancy at this time. Broadcast and multicast transfers are possible through the use of logical addressing. It should be noted, though, that this introduces inefficiency into the network and the high probability that all frames will not be transferred due to paths already being in use.

As one module can connect to every other module via the transparent Fabric, modules executing applications which need to communicate do not need to be in the same rack since throughput and latency will be predominantly the same. Obviously there are advantages in local transfers (i.e. using a single switch) since it minimises possible path conflicts and inter-switch bottlenecks. The size of the network/Fabric is governed by the number and size of the switches. Being based on a commercial standard, component developments are likely to be available in large quantities, allowing development costs to be amortised over many devices, and competition will help to drive component prices down. It is noted that this would not include the maturation of components for the military environment.

Disadvantages of the system include the presence of active elements in the data path, causing unreliability of individual paths, although there are few single point failures The fact that a single network is used also lives concern over segregation for security. This would require careful and explicit consideration. Translating the FC standards to a usable system would require the upper layers and intended system architecture to be specified.

2.6.5 Performance

As stated above, any module can connect to any other module, but one transfer does not impact another if the same resources (i.e. transmit and receive paths) are not involved. In particular, a module can simultaneously transmit and receive transfers involving different modules. Generic estimates of message transfer performance per link can be made of 217/244 Mbit/s (global/local) and a data streaming performance of 835 Mbit/s. By making the resources dedicated between two modules, this level of performance will always be available. Multiplexing via the switch between different connections will reduce this capability due to contention over resources, but would allow more general use. Multiple, independent links in the system, many operating in parallel at any one time, allow enormous potential throughput. A module input/output is limited to 1.6 Gbit/sec (half input, half output) through a single interface.

2.6.6 Implementation

The FC links between modules in a rack would be implemented using all optical waveguide backplane with the switches implemented in separate modules. The use of' an optical backplane is expected to be common to all network approaches addressed in this study. The backplane in this case will be fairly simple and common to all racks, comprising point-to-point links from the processing module slots to the switch module slots.

Links between racks would be by fibre optic cable. Multiway optical connectors would be used between the backplane and the modules (both processing and switches) as well as on the rack to interface to the inter-rack fibres. Each module will have a single interface (probably two for redundancy, but this has not been further considered). The exact number of switches in a system and their connection to the modules has not been considered in detail, partly because no information concerning switch implementations has been obtained. In the performance estimates it has been assumed that a maximum of two switches would be traversed by a single transfer, but this is an arbitrary limit determined as a result of the overhead of multiple links being employed. The inter-rack connections for the system are fairly minimal and simple point-to-point links.

A single interface (or a dual redundant interface) is considered to require a small board area (~4%). The single (dual) interface also minimises the fibre optic connection problem between the processing modules and the back-plane, although a large number of connections would be required for the switch module to the back-plane.

The network would operate at 1.0625 Gbaud using 8B/10B coding to provide a maximum peak- data rate of 850 Mbit/sec.

2.6.7 Risk

The main risk items for an FC based system is the need for multiway monomode military connectors, electro-optic device integration and LASERs stable in an avionic environment. All of these are common to all the network approaches.

The development of switch fabric components for FC will be driven by commercial system requirements. Minimal development would be required to adapt these for avionic environments. However, it is noted that the use of FC within commercial systems will concentrate on the use of Class 1 connections (dedicated resources) with connections maintained for very long times (hours) which would not be suitable for the message passing Class 3 datagrams envisaged to form an important part of the avionic implementation.

2.6.8 Growth

Growth of the FC concept must be considered from the viewpoint of encompassing new technology and allowing implemented systems to increase their functionality. Fibre Channel is identified to operate at a variety of speeds already and further increases in speed can be expected to be supported by commercial systems. Such increases would be simply incorporated by the replacement of processor module interfaces and switch modules. Larger, more capable switch implementations allowing greater connectivity can also be expected. Reconnection of modules (e.g. replacing two switch modules with a single one) would be accomplished through the introduction of a new backplane.

Growth within an existing system would be achieved through spare slots within the rack connected to spare switch ports and spare slots for additional switch modules. This would be easy to define during system design for P^3I.

2.6.9 Standards

The standards are promoted by the Fibre Channel association who have a huge www site. http://www.fibrechannel.com. A listing of this web site was circulated to HS WG of ASSC.

2.7 Giga-bit Ethernet

2.7.1 Background

A comprehensive up to date description of status is available at http://www.gigabit-ethernet.org/.

Ethernet has always suffered a bad technical press since it has a spectacular and easily understood failure mode. In reality this is easily avoided by judicious deployment. It is a fact that Ethernet (100 Base-T) accounts for 83% of all installed LAN connections some 150 Million terminations. Current deliveries reinforce this dominance. In the commercial world Ethernet has emerged as the inexpensive trouble free, easy to deploy LAN. Token ring is the nearest rival with FDDI and others. The dominance of traditional Ethernet results in desire to exploit the name on ‘go faster’ protocols. 100 Mbit/s published in ‘95, 1Gbit/s to be balloted ‘98, 10 Gbit/s is in gestation, declared to be feasible and scheduled for circa 2002.

2.7.2 Description

Gigabit-Ethernet is being standardised by IEEE 802.3z and has a compatible frame format to the existing 10 Base-T and 100 Base-T. The market goal is ‘backbone’ transmission between Ethernet bridges. The essential CSMA/CD medium allocation strategy is very well known. Gigabit Ethernet is a star where the contention takes place at a point. Each user fires packets to the star which broadcasts the packet along all other links. 100 Base-T proved the concept. There are several physical layers. The default physical layer is a straight importation of the Fibre Channel 1.24 Bits FO Physical layer. There are options for local systems using multi-mode fibre, others for single mode fibre offering several kilometres. There are even copper links capable of 1.24 Gbit/s over 25 metres using 150 ohm twinax cable. The 1.24G results from 1G data transfer with 8B10B encoding.

2.7.3 Performance

Performance is still to be determined. The picture is becoming complicated by recent proprietary initiatives that avoid contention by buffering. Performance is also a function of packet size. Small packets are inefficient. With contention, throughputs of 700 Mbit/s are deemed feasible by modelling. With ‘buffered distributor’ then throughput approaches offered traffic even at 100% load (1G of offered traffic). Presumably throughput is maintained at the expense of some latency and jitter.

2.7.4 Implementation

The physical layer standards are imported from proven initiatives elsewhere. This means that proven Fibre Channel components and modulation are used to build the star. The function of the switch in the centre is complicated by several proprietary initiatives. In practice Ethernet bridges and routers evolved the same way, i.e. they conform to ‘the standard’ externally but the internal resolution is proprietary.

2.7.5 Risk

This depends on perspective. The market sees Ethernet evolution as ‘least risk’ solution. Their investment is with Ethernet packet format and bridging technique. The preservation and evolution of the existing packet format reduces total risk by a huge amount. The physical technology ie fibre, OE modulation is proven. The commercial pressure to succeed is enormous. The market lifetime of products is likely to be long (decade) compared to some other initiatives.

2.8 Universal Serial Bus (USB)

Two major new interface technologies are appearing for desktop personal computers which due to their very low cost and widespread industry support may well appear in future avionics systems applications. These interfaces are the Universal Serial Bus (USB) and IEEE 1394 otherwise known as "Firewire".

The USB is intended to replace the older 'RS' series of serial interface standards such as RS232,422,485 and the 'Centronics' parallel ports for serving low speed peripherals such as mouse, keyboard, scanner and printing devices. The IEEE 1394 interface is intended to replace SCSI interfaces for higher speed devices such as hard disks, CD-ROM's, digital video, etc.

Both these new interface technologies feature common characteristics to support modern 'hot insertion and removal' and 'plug and play' capabilities. The USB features Asynchronous and Isochronous modes of operation over a simple common cabling scheme which also provides the capability of distributing power to peripherals within the standard cable. The IEE 1394 interface is covered in section 2.5. The main features of USB are as follows.

USB provides support for up to 127 logical devices connected via a tiered star topology (Christmas tree) using multi-way 'hubs' of up to six tiers. The system supports both 1.5 and 12Mbit/s operation with control being dedicated to one USB host, normally in the host computer/PC. 'Hot' insertion and removal of peripheral devices is managed by automatic re-configuration software/firmware within the host.

The physical layer consists of 2-wire differential signalling along a twisted pair at CMOS 3.3V levels. Transmitters are controlled to have rise and fall times between 4 to 20 nsec to minimise RFI and signal skew. The twisted pair is combined with 2 power conductors to form standardised interconnection cables consisting of shielded cable assemblies of up to 5 metre lengths terminated in standardised 4 pin shielded (blind mating) connectors at each end. The twisted pair has a line impedance of 90 Ohms and the power conductors are rated for the supply of 5 Volts at up to 500mA.

The USB system topology consists of three basic elements:-

(1) The Host Controller which is responsible for centralised command and control of the network being most typically found in the host computer PC/workstation.

(2) The Hub which is a router/repeater device for providing the connectivity for the peripheral devices. Every hub has one "upstream port" and one or more "downstream ports". In personal computer applications keyboards and monitors often provide the USB hub function and typically support at least *four* downstream ports.

(3) The Peripheral Interface is the actual interface providing USB support for the attached devices and they come in two forms, the Low-Speed Peripheral Interface (LSPI) for 1.5Mbit/s operation and the Full-Speed Peripheral Interface (FSPI) for 12Mbit/s operation.

The USB system supports a number of different data transfer types to suit the varying real-time requirements of the attached peripherals. These are primarily in the form of data or control exchanges between the host and peripheral devices either as uni or bi-directional transfers. Communication is set up between the host and either single or multiple endpoints and such associations in USB terminology are called 'pipes'.

There are four main transfer modes consisting of *Isochronous* mode for delivery of real time messages at a guaranteed rate, *Interrupt* mode for delivery of data no slower than a peripheral device specifies, *Control* mode for configuration of peripheral devices when they are first attached to a USB network and *bulk* mode for delivery of large blocks of data requiring high bandwidth but at infrequent intervals.

Transfers are controlled by *token* packets which are only issued by the USB host to allow for *in* , *out* and *setup* type transactions. The data field may range from 0 to 1023 bytes which must be an integral number of bytes. The data bits within each byte are transmitted LSB first. *Handshake* packets are transmitted to report the status of data transactions and can return values indicating successful reception of data, flow, control and stall conditions. *Start of frame* packets are issued by the USB host at a nominal rate of 1msec and consist of a packet ID followed by a 11 bit frame number field. Both the token, data and start-of-frame packets are protected with 5 bit CRC fields.

In many ways the centralised control philosophy of the USB together with the transfer modes described above resemble the well accepted command/response philosophy of the Mil-Std-1553B (Def-Stan 00-18 (Part 2)) avionics data bus. This would make the USB, as a modern COTS derivative, attractive for potential avionics system use. However at the time of writing the USB is only just becoming widespread in the personal computer world and there are no known avionic applications (although some avionics related research projects feature its use). Therefore before use in an avionics environment uncertainties about the robustness of USB, particularly for EMC, as well as the limitation of the tiered star topology would have to be evaluated for avionics system use. However the shear size of the market for USB will make components very cheap and readily available which will in turn encourage its use for avionics system applications and hence its inclusion in this guide.

2.9 ITU-T Rec. X.34 (1996) - Data networks and open system communications

ITU-T Rec. X.34 (1996) covers data networks and open system communications for public data networks – Interfaces; and provides recommendations for access to packet-switched data transmission services via B-ISDN

2.10 ITU-T Rec. X.96 (1996) - Data networks and open system communications

ITU-T Rec. X.96 (1996) covers data networks and open system communications for public data networks – Interfaces and provides recommendations for procedures for the exchange of control information and user data between a Facsimile Packet

Assembly/Disassembly (FPAD) facility and a packet mode Data Terminal Equipment (DTE) or another FPAD

3 STRENGTHS AND WEAKNESSES

The strengths and weaknesses of standard bus must be considered before any decision can be made regarding whether the performance of a standard meets the requirements for the system. This may not be simply the peak performance of the standard but may consider whether the operation of the standard is adequate to enable the system to perform within given margins.

Other criteria which may need to be considered include: latency determination, interface power dissipation, interface size, data transmission system weight, system operation (periodic/event driven), fault tolerance, failure modes, growth capability, cost, VLSI interface availability, previous experience, EMC performance and fibre v copper considerations.

The avionic environment has several well known environmental requirements that probably exceed any civil specification. A fibre optic and its terminations specified by BT for FTTC might not endure the shock, vibration, temperature excursions, etc. of a fighter aircraft. Moreover the installation requirements are very different, e.g. bulkheads, flammable atmospheres, resilience to battle damage. The majority of these issues are crucial to selecting fit for purpose economic solutions.

4 COMPARATIVE PERFORMANCE

This section compares and contrasts two aspects of the four standards. A comparison of performance is provided. A comparison of the other assessment criteria must be made by reference to the Section 2 where each of the standards is described in turn. Other issues such as initialisation, delays and redundancy are not reported because the processes are not adequately described in the standards, or are specific to the application in which the data network is designed to operate.

The current and projected data rates of' the four standards are given in the table below.

Standard	Current Maximum Data Rate	Projected Data Rate
Fibre Channel	850 Mbit/s (1062.5 Mbaud)	2 Gbit/s
ATM/SDH	622 Mbit/s	2.4 Gbit/s
SCI	1 Gbyte/s (parallel) 1 Gbit/s (serial) (1.25 Gbaud)	2 Gbit/s
Gigabit Ethernet	1 Gbit/s	10 Gbit/s

Table 1 Projected data rates

It can be seen from the table that all the standards are moving towards the same overall performance capability. However, peak data rate is in itself not an adequate measure of a standards' performance. It is not simply a matter of choosing the fastest one, it is necessary to identify its capability when a particular data flow scenario is applied.

5 APPLICATIONS

Interface standard	Application
Asynchronous Transfer Mode (ATM)	Global CO telecommunications
Scalable Coherent Interface (SCI)	Distributed computing, work-station clusters
Fibre Channel (FC)	workstation, mainframe, peripheral interconnect
SerialExpress (P.2100)	PC, workstation, ‘plug and play’ peripherals
Giga-bit Ethernet	LAN-LAN backbone especially Ethernet LAN

Table 2 Indications of interface standard application to current systems

6 CONCLUSIONS

6.1 General

A brief paper exercise to consider the merits of these standards can do little more than conclude that none are likely to be directly ‘fit for purpose’ in an avionics application. To a first approximation all the protocols discussed could probably meet the general technical requirements of an avionics Network. This has to be qualified since for example, the way ATM as used is unsuitable for military avionics applications. The objective is that military versions of ATM should be the same as used in civil applications. The consequences of a specific project departing from the specification should be realised. Other protocols are likely to have attributes that are prejudicial to a military mission. Identifying and work-arounds for any such deficiency is not a simple matter.

Whilst the protocol could be slimmed or otherwise amended all these contenders would be more or less deficient in respect of their common physical components. In other words the protocol might be adequate, however the civilian components are probably deficient.

The notion of a unified high speed network, serving avionics in the manner that US MIL-STD-1553B has done, is academically feasible but unlikely in the near future. The probability is that existing certified solutions (using US MIL-STD-1553B) will continue and new functions are added using one of the civil standards under consideration here. The final decision should be made per project. It is sensible to do the hands on evaluation in advance of a project otherwise proprietary solutions might be expedient or the design compromised by insufficient time for full evaluation.

The cost of ownership for any of these protocols requires consideration for the actual applications. Each commercial initiative is designed to serve a particular marketplace. Within that marketplace the communications resource and the application are mutually dependent. It is probable the real cost of an avionics system will be as much influenced by the second order COTS opportunities in applications as with the interconnect.

7 GLOSSARY

7.1 Definitions

4B5B Encoding	An encoding method whereby four information data bits are encoded into five-bit symbols.
ADSL	Asynchronous Digital Subscriber Loop (ADSL)
ATM	Asynchronous Transfer Mode
Bandwidth	The carrying capacity of a network, usually measured in bits per second for digital circuits, or Hertz for analogue circuits.
Connection oriented service (CO)	One in which the parameters for a series of transfers are given before any data is passed and thereafter the user provides data only.
Connectionless service (CNLS)	One in which the data and parameters for the transfer are provided by the user for each transfer.
COTS	Commercial ‘off the shelf’
Data rate	The speed at which bits are transmitted and received, usually measured in bits per second.
Deterministic	Property of an item to which the future behaviour can be predicted precisely.
Efficiency	The percentage ratio of the number of data bits to the total number of transmission bits (including overhead).
Fault Tolerance	Capability of the system to endure component errors and/or failures without causing total system failure. Actions range from ignoring it, to retrying action, to complex actions of fault isolation and then taking positive action to continue operation without the failed component.
FDDI	Fibre Distributed Data Interface
Flow control	The ability to limit the rate at which a terminal can transmit data.
FTTC	Fibre to the Kerb. The delivery of fibre connections to end users.
Full duplex communications	A simultaneous, two-way communications path.
Half duplex communications	An alternating transmission path, two ways, but only one direction at a time.
IEC	International Electro-technical Commission
IEEE	
ISO	International Standards Organisation
Latency	A measure of time delay.
Manchester II encoding	A way of encoding to get a zero-DC binary waveform. In this encoding scheme, half of the bit interval is transmitted with a positive signal and the other half is transmitted with a negative signal.
Overhead	Non-data bits or characters necessary for transmission, error detection or for use by protocol.

Packet	A block of data, including addressing, routing and numbering information. Very long messages may be segmented into packets for transmission over the data transmission path.
Protocol	A set of related rules describing specific processes or activities.
SDH	Synchronous Digital Hierarchy
Signalling rate	The speed at which data bits and overhead bits are transmitted sometimes referred to as Baud rate.
Simplex communications	A one direction communications path.
Throughput	A measure of the effective network transmission speed, it is the net bandwidth of a network.

8 REFERENCES

8.1 Document references

1. Society of Automotive Engineers (SAE), Handbook of System Data Communications (AIR 4271), November 1989.
2. Relevance of commercial standards, June 1994 - Report by GEC-Marconi Avionics, DERA, Farnborough, and GKN Westland Helicopters Ltd (Document ref: ASSC/120/3/164.
3. Transfer facilities report for high performance networks - Report by the Mitre Corporation (Document reference ASSC/120/3/158.
4. IEEE, Scalable Coherent Interface (SCI) Logical, Physical, and Cache Coherence specifications, IEEE standard 1596, 1992.
5. ANSI, Fibre Channel Physical and Signalling Interface revision 4.0, working draft of ANSI committee X3T9.
6. ATM Forum, ATM User-to-Network Interface Specification Version 3.0, publication pending by Prentice Hall.

8.2 Web site references

	Title	**URL**
1.	Fibre channel association	http://www.fibrechannel.com
2.	Fibre Channel (FC) projects	http://www.dpt.com/t11/fc.html
3.	FC AE	http://www.bardstown.com/~fcae/
4.	Fibre Channel (FC) Association server	http://www.amdahl.com/ext/CARP/FCA/F html
5.	Gigabit Ethernet	http://www.gigabit-ethernet.org
6.	SCIzzL Overview	http://www.scizzl.com/
7.	Serial Express	http://www.SCIzzL.com/P2100/index.html
8.	ISO/IEC JTC1 SC21: Open systems interconnection, data management and open distributed processing	http://www.iso.ch/meme/JTC1SC21.html
9.	ISO/IEC JTC1 SC25: Interconnection of information technology equipment	http://www.iso.ch/meme/JTC1SC25.html
10.	The ATM Forum Home Page	http://www.atmforum.com/atmforum/