Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

100TH T11 PLENARY SPECIAL EDITION

	Page
Synopsis	2
Protocol Engine Meeting Jan 15 1988	3
History	4
Applications	5
Performance	6
Perspective	8
X3T9.3 IPI Meeting Feb 22–23 1988	8
Fiber IPI	8
X3T9.5 FDDI Meeting Feb 24–25 1988	9
Fiber Channel SG June 2–3 1988	9
Rationale	11
Components	12
Requirements	13
Technology	14
Perspective	15
X3T9.3 IPI Meeting June 20–21 1988	15
Fiber Channel	16
Fiber Channel SG Aug 11–12 1988	17
Protocol Engine	18
Missing the Boat	20
Optical Primer	21
Lasers	22
SCSI Fiber Extender	27
DQDB	28
Parallel Connectors	29
MacroLAN	30
HP-FL	32
Switches	33
Protocol	34
Fiber Channel SG Sept 26–27 1988	34
Packet Switching	35
Types of Switching	36
Objectives	37
Requirements	38
Definitions	39
Application Area	39
Environment	39
Cost/Distance	40
Fault Tolerance	40
Performance	40
Optical Extender	40
SD3 Status	41
What About FDDI?	42
Fiber Channel WG Dec 1–2 1988	42
Office Environment	43
Topology	44
Principles of Operation	45
Application	46
HSC	47
Principles	48
Working Documents	50
Acronym Glossary	50

Congratulations T11 and FCIA

ENDL prepared this special edition of the Happenings to support the FCIA (Fibre Channel Industry Association) in celebrating the 100th T11 Plenary on Thursday August 5, 2010. A PDF copy of this Special Edition is available at: http://www.endl-letter.com/se100t11.pdf.

Permission to distribute this Special Edition is granted to FCIA Member companies.

14426 Black Walnut Court, Saratoga, California 95070 408-867-6630 endlcom@acm.org

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

SYNOPSIS

PROTOCOL ENGINE MEETING JANUARY 15 1988

The pressures behind the ambition to streamline networks proved useful ammunition in the arguments used to justify pursuing a storage-oriented solution.

X3T9.3 IPI MEETING FEBRUARY 22–23 1988

The procedural underpinnings for Fibre Channel were established with approval of the IPI-3 project proposal.

X3T9.5 FDDI MEETING FEBRUARY 24–25 1988

The growing tendency to build intelligent peripherals was noted.

FIBER CHANNEL STUDY GROUP JUNE 2–3 1988

The IPI-3 project blossomed into (no fair reading ahead) Fiber Channel. Mountains of preparatory transparencies were presented and reviewed (stone knives and bear skins only, no PowerPoint, and video projectors too big fit in car's trunk, never mind a suitcase). By the end of the gathering, the broad FC outlines were more or less understood.

X3T9.3 IPI MEETING JUNE 20–21 1988

Requirements were reviewed, and addressibility was limited to 256 devices. This would eventually affect Arbitrated Loop, but Fabric Fibre Channel would blow through this like a twister through Oklahoma.

FIBER CHANNEL STUDY GROUP AUGUST 11–12 1988

Signaling issues and education were Job One. A few adventures into higher layers netted ideas that would later worm their way into the FC-x levels.

FIBER CHANNEL STUDY GROUP SEPTEMBER 26–27 1988

Switching concepts made their first serious visit to the group. Requirements were also updated across the board.

FIBER CHANNEL WORKING GROUP DECEMBER 1–2 1988

During this everything-but-the-kitchen-sink gathering several of the building blocks that are recognizable to this day made their first appearance.

IN THE BEGINNING

Fibre Channel was not a sudden concept, and the circumstances that justified such a project were not new. All of the solutions being considered were based on networking. The idea that a storage interface could be networked was anathema to both network and storage bigots alike.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

1988 was the critical year, the first study group was held in June, and by year end the project was off and running. By way of prelude, extracts from meetings earlier in the year have been included to paint a picture of the circumstances which preceded the study group.

It was politically unacceptable to suggest Fibre Channel was a network, any hint that it could be so met instant opposition from within (X3T9.5) and without (IEEE) the X3T9.3 committee. It was almost a decade before the term Storage Area Network was introduced for Fibre Channel. Any why the odd spelling might you ask? The journey to eventual project approval in 1989 was both difficult and conditional e.g. one of the criteria for project approval was that 1988's spelling of Fiber become Fibre.

For those of you who were there, this should revive memories. Those who were not should gain an appreciation of the pioneering done by those early participants.

I. Dal Allan
endlcom@acm.org

Ralph O. Weber
roweber@ieee.org

PROTOCOL ENGINE MEETING JANUARY 15 1988

Larry Green, formerly of ACC, founded Protocol Engines to bring to market a chip set which Greg Chesson (Silicon Graphics) has been working on to automate Layers 3 and 4 of the ISO model.

A TAB (Technical Advisory Board) is being formed to nurture the project. TAB membership runs $25,000/year. Meetings are normally closed but this one was free and open to all as a way of encouraging companies to join. The size of the crowd was indicative of the interest; over half had flown in especially for the meeting, held at the Fess Parker Red Lion Hotel in Santa Barbara.

Company	Name	Company	Name
AMD	F. Firoozmand	British Telecom	T. Boston
Ameritech	G. Nelson	Brooktree	M. Eichen
Apollo	B. McClure		J. Teza
AWA Australia	R. Halgren	ComDesign	S. Beisner
	C. Tran	Consultant	C. Bender
BNR	P. Chen		

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Company	Name	Company	Name
Corp Open Systems	M. Cain	PCO	M. Hartmann
	T. Manakas	Proteon	H. Salwen
CTA	L. Neitzel	Protocol Engines	R. Dolan
Delco	J. Chelini		S. Golbuff
EDS	D. Carros		L. Green
ENDL	D. Allan	Reetz & Fox	R. Fox
FiberCom	T. Chan	Silicon Graphics	F. Baskett
GEC Research	P. Clarke		I. Weidhof
Interphase	M. Cope	StanaTek	J. Hamstra
	J. Foster	STC Tech	K. Caves
	M. Lease	Stellar Comp	R. Gurwitz
Martin Marietta	R. Hewitt	Oracle	D. Scherbath
	E. Hotard	Systems Synthesis	D. Brain
Merrill Lynch	M. Mahoney	Texas Instruments	J. Carlo
National Semi	R. Calvo	U of Virginia	A. Weaver
Naval Weapons Ctr	D. Marlow	Ultra	K. Bossange
	K. O'Donoghue	Unisys	S. Andersen
Northrop	M. Cohn		F. Ross
	J. Walker	Wavefront Tech	L. Barels
PCM Svcs	S. Cooper	Xerox PARC	J. Larson

A limit of 12 companies has been set on the TAB, and at meeting time, four had signed up: Apollo, Interphase, Martin Marietta, and Silicon Graphics. NOTE: Unisys and Xerox joined subsequent to the meeting.

When the four TAB members described the reasons they had joined, all but Mike Cope (Interphase) spoke of generic benefits. Mike sees communications as a diversification strategy - Interphase derives most of its income from disk controllers. The time is right to apply technology learned on other peripherals to the "tremendous potential of high performance communication products". The fact that Mike has no installed base to worry about is added incentive to ignore the past and concentrate on futures.

History

Greg summarized deficiencies of current protocols and described how gateways between different networks on a coterie of media will be the norm, not the exception. Unfortunately, the efficiency of existing protocols is so low the bandwidth of a backbone such as FDDI will be miniscule.

As an example of why setting standards became one of the primary goals of Protocol Engine, he related the story of how a number of proposals for ISDN were solicited. Siemens had the best, it was two years ahead of every other country.

Every other proposal was a paper study but adopting the Siemens proposal was unacceptable as it would give Germany an edge in developing products.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The result was an emasculated ISDN which put everybody on an equal footing. The moral of the story is that it is wise not to get too far ahead in technology or the politics of the standard bodies will zap you.

When it comes to effort spent on the Protocol Engine, the transport protocol is actually the easiest part of the job. The hardest part is ensuring fast operation of the supporting hardware and software necessary to get bits in and out. About 80% of the effort is in ensuring that implicit buffering and I/O paths are integrated in order to achieve high performance.

Applications

Larry Barels (Wavefront Technologies) specializes in graphics software for the entertainment, aerospace, and super computer markets. His goal is to provide three dimensional photorealism on a screen.

Unix will continue to be the basis for all graphics growth - it is provided on over 80% of all work stations and it is the most popular operating system on Cray and the minisupers. FDDI is a given for graphics applications, and it is essential that the maximum speed be obtained in order to obtain the kind of display speed that is needed.

Note: It seems strange that Larry talked of FDDI at 100 Mbs when it seems an even more likely transfer rate for the kind of performance he described is 8 times faster at 100 MBs.

Ilana Weidhof (Silicon Graphics) has just returned from a European survey, done to get a handle on prospect and customer expectations from LANs. The European business scene is booming, and the infrastructure is falling apart. Mergers and acquisitions are occurring at a hectic pace, Mercedes Benz alone bought 30 companies in 1987. Stiff competitive pressures and an affordable dollar due to exchange rate adjustments is causing the seams of industry to render under massive changes.

At Mercedes Benz, the computer center once had 50 stations attached, but it now has hundreds of users, and next year it expects to have thousands. The mainframes are being choked by the demand for time. The triangle of activity is between computing, data storage, and visualization. IBM is looked upon as the data storage facility for the compute center (typically a Cray or some other super computer). Visualization is what people want most, and growth in graphic representation appears limitless.

One thing Ilana wanted to find out was what users are looking for in the way of OSI (Open Systems Interconnect) protocols. Rather than answers, she heard back questions about what was going on in FDDI. The common theme amongst all was performance, the biggest problem is getting data in and out of the IBM data storage system quickly. The important factors that came through universally include:

- They want results
- They want to share them
- Applications are what matters
- Integrated environments (e.g. DEC and IBM)
- They need all these abilities now, not later

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

David Brain (Systems Synthesis) described the objective he has at Maxwell Publishing in England, of integrating and tying together the special purpose printing equipment installed at multiple sites. Each has its own closed system interfaces. It is impossible to move the required volume of data between sites so that Maxwell can take advantage of the capabilities of each special purpose piece of equipment.

At present speeds, it is cheaper and faster to put the data on a motorbike than it is to transfer it through communication links. He is looking forward to FDDI, and efficient transfer structures that would allow him to link the various sites into a realistic network.

Marc Cohn (Northrop) described SAFENET (Survivable, Adaptable Fiberoptic NETwork). The SA in the acronym originally stood for Shipboard Applications but it has moved from the Navy to cover all forces. SAFENET differs from any other defense effort because the primary goal is utilizing commercial standards to come up with a militarized open system architecture.

XTP, the Protocol Engine transfer protocol, and the Protocol Engine chip(s) were described by Marc as critical to SAFENET. The biggest advantage will be interoperability. Tying the pieces of a network together is a nightmare at present because so many subcontractors do different parts of the work. A dedicated contractor has to be paid to ensure interoperation between them all.

ASC X3S3 is responsible for Layers 3 and 4 of the OSI model. Chairman Lyman Chapin (Data General) is sympathetic to considering XTP in committee. NOTE: Subsequent to the meeting, the SAFENET standards group met and adopted XTP. It was a unanimous vote from a group composed primarily of XTP beneficiaries - the users. This could influence X3S3.

Performance

Forest Baskett (Silicon Graphics) spoke to the scandalous nature of present protocols which are so poor in performance that Ethernet can achieve less than 1 Mbs throughput out of a 10 Mbs bandwidth. The scaling is not linear; on an 80 Mbs Proteon installation, the throughput is only 6.3 Mbs.

Howard Salwen (Proteon) added that the bottleneck occurs trying to get data in and out of the system. Cray is able to run the protocols at 110 Mbs internally but a typical file transfer takes one second due to the 8 interrupts per packet.

Howard remarked privately that Martin Marietta, using a homebrew lightweight protocol, had achieved 37 Mbs on Proteon. As the theoretical tops is 45 Mbs given the design constraints under which it was operating, this is over 80% efficient. It is a clear illustration of lightweight protocol advantages.

Greg described the typical design of a communications controller.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

If Greg's portrayal is accurate, communication controllers appear to be at least five years, and perhaps as much as a decade, behind disk controller technology.

Alf Weaver (Center For Communication Research, University of Virginia) has a somewhat different dilemma to those who want to move more data; he is working under contract on a realtime NASA application with small chunks of data which require very low latency response.

To evaluate the existing protocols, he used Intel iSBC 80286s with a token bus card that has its own 80186 processor to generate data, and measured performance using the latest generation optimized software he could find.

Alf's purpose was to figure out how long it took to get from the transmitter to the receiver using the OSI model. It took 9 msec to get started at the Data Link Layer. Net throughput came to 65 KBs out of 10 Mbs bandwidth. The maximum rate achieved with tiny messages was 160/second, and these messages did not require acknowledgment.

Using the same hardware, but a Transport Layer interface with messages being acknowledged to ensure data integrity, it took 20 msec to transfer initial segments up to 1K, and 5-15 msec for later segments of the same message. The given overhead for end-end times was 100 msec.

Performance topped out at 72 Kbs with 10K transfers. The profile was a sawtooth curve, as each new segment caused a hit that created jaggies of inefficiency. The best that could be achieved was 55 messages/second at 16 bytes down to a low of only 10 messages/second with 16 KB packets.

The LLC (Link Layer Control) software interface to the MAC hardware is poorly designed in all systems. It unrealistic in realtime applications to be forced to wait for 100 msec of delay before timing out on a lost message.

The Space Station will have three networks, one of which will be connection oriented with a low BER (Bit Error Rate) and the other two will be datagram services with a less stringent BER. The requirement for interoperability via the OSI protocol causes inefficiencies that will have significant impact on performance. OSI was designed for text transmission between companies, not for real time systems applications.

The best case routing of space to ground transmission passes through approximately 20 layers from transmitter to receiver. The NIU (Network Interface Unit) is a Honeywell contract, and NIU is an entity separate from the host. None of the elements to be integrated by NASA have known performance, and the subcontractors do not have to share information during development.

In the Space Station system, ignorance is not bliss, it is low performance.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Perspective

Two hours of listening to Greg does not an expert make, but there are XTP/Protocol Engine features which seem to parallel the functionality provided in IPI e.g. XTP can multiplex by signifying the end of known data during transmittal. This and other similarities need to be investigated deeper, but it suggests there is potential for IPI sitting on top of the Protocol Engine on an FDDI or other links.

There is a flip side advantage also. If a Protocol Engine chip set was designed for the higher IPI speeds it would be possible to put communications behind the Protocol Engine on a Physical Interface. Plans to originate a FIPI (Fiber IPI) are underway, so this could be extremely attractive to those wanting to make image transfers using the same software as the rest of their network.

X3T9.3 IPI MEETING FEBRUARY 22–23 1988

Some who signed in mostly attended SCSI to listen to the connector debates.

Company	Name	Company	Name
3M	J. Filia	GP Associates	R. Grow
Amdahl	E. Cardinal		R. Perloff
Aptec	R. Phillips	Hewlett Packard	D. McIntyre
Bull	P. Schreck	Honeywell Bull	D. Barney
CDC	W. Sanderson	IBM	H. Brandt
CDC MPI	T. Leland		E. Calkins
Consultant	M. Baird		D. Roberts
Convex	G. Steagher	IBM Research	P. Zafiropulo
Cray Research	S. Johnson	Integrated Photonics	W. St. John
	J. Renwick	Los Alamos Lab	D. Tolmie
DEC	G. Robinson	Maxtor	D. Banerje
Eastman Kodak	R. Derr	Network Systems	K. Drewlo
ENDL	D. Allan	Prime	M. Fitzpatrick
Fujitsu	M. Moteji	Scientific Comp Sys	D. Deal
Fujitsu America	R. Driscal	Simulex	R. Morris
	J. Luttrull	StorageTek	D. Appleyard
Gould CSD	E. Balogh	T&B	P. Phillips
	D. Healy		H. Waltersdorf
	M. Klasey	Ultra	N. Perdue

Fiber IPI

The Enhanced IPI may not have received committee approval to go forward to X3T9, but the Fiber IPI SD3 fared better. There was some discussion about its failure to mention anything about transfer rates and opposition to mentioning X3.129 and Enhanced IPI as companion documents.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Despite having sat in the IPI committee for several months the HSC-oriented members have little or no understanding of the concept that levels provide technology independence. Even after it was explained, there was a general feeling of disbelief, but come time to vote, it was forwarded unanimously.

X3T9.5 FDDI MEETING FEBRUARY 24–25 1988

Ray Alderman (Matrix) was far better, describing the problems facing board makers in the future. Ray sees the market evolving from boards into subsystems. Packaging is a critical skill that must be mastered. Proper grounding and cable shielding is atypical. Bus clock speeds are increasing, yet there is little attention given to interconnect capacitance.

> "With smarter peripherals and increasingly complex protocols, will we be running on an optical fiber link in the future?"

FIBER CHANNEL STUDY GROUP JUNE 2–3 1988

The meeting was hosted by AMD at the Holiday Inn in Sunnyvale.

Company	Name	Company	Name
Adaptec	R. Snively	Fujitsu America	J. Luttrull
AMD	J. Jaeger	Gazelle μCircuits	C. Popat
	J. Kubinec	Gigabit Logic	J. Kemps
	P. Scott	Hewlett Packard	D. Hanson
	R. Sterner		S. Joiner
AMP	C. Brill		I. Ohel
	G. Moore	IBM	G. Heiling
AT&T	R. Knerr		H. Truestedt
	J. Morris	Integrated Photonics	M. Pugh
	R. Travieso	Lockheed	J. Van Saders
Avantek	G. LaBelle	Los Alamos Labs	D. Tolmie
BT&D Tech	J. Lowell	Lytel	P. Clarke
Canstar	K. Lue Shing	MMD	P. Ralston
CDC	W. Sanderson	NCR	J. Lohmeyer
CDC Canada	R. Cummings	P&B/Siemens	S. Van Doorn
Data General	S. Solomon	Raycom Systems	M. Nakamori
ENDL	D. Allan	Unisys	B. Bergey
Fujitsu America	B. Driscal		P. Dougherty

Although originally chartered as an IPI Physical Interface alternative, the consensus was that it should be expanded to become the FC (Fiber Channel), with the stated intention of supporting multiple command sets.

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The meeting opened with a summary of its purpose:

Gather Data on Optical Transmission		
Components Connectors	Coding Schemes Generic Protocols	Transmission Media Error Considerations
Establish Functional Requirements		
Cost Speed Distance	Performance Applications Environments	Fault Tolerance

The purpose of the project was summarized as:

Physical Interface
Define Level 0 transmission components Define Level 1 protocol and control mechanism
Logical Interface
Support IPI-2 and IPI-3 Command Sets Support SCSI and SCSI-2 Command Sets

The intent of the project is to ensure that the software investment in SCSI and IPI is retained.

A number of starting considerations were listed:

Protocol	Components	Media
Datagram Connection	LEDs Lasers	plastic multimode single mode

Del Hanson (Hewlett Packard) queried what the differences were between the FC and a LAN (Local Area Network), because he saw them as being more similar than different. True, many of the problems and solutions will be the same as a LAN, WAN (Wide Area Network) or MAN (Metropolitan Area Network).

The primary difference is that the FC is not intended to conform with the OSI (Open System Interconnect) model but operate in a known, specified, and limited domain.

One obvious difference is limited addressability: LANs work with 48 bit addresses but there is seen to be no need for more than 32 units on the FC. Another difference is the requirement to support a broad range of speeds and a variety of transmission mediums. The emphasis in small system applications is low cost, yet the same software has to be able to execute on a very high speed link between mainframes, superminis and supercomputers.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Rationale

The companies present were split fairly equally between users and vendors:

Users	
Canstar	Currently selling a 100 Mbs network and investigating FC to see if its anticipated higher performance could be used in a network
CDC	Wants a low latency interface which can go a further distance than copper, and be modular from 100-200 MBs
Data General	Looking for a faster, more fault tolerant interface than FDDI
Gazelle µCircuits	Exploring standards with abilities in the GHz range
IBM	Leave copper behind
Lockheed	Needs something along the lines of HSC in fiber
Los Alamos Labs	Packaging and distance improvements over copper HSC
NCR	Interested in using SCSI on fiber across all systems
Raycom Systems	Has a goal to replace copper with fiber for reduced Electro Magnetic and Radio Frequency emissions
Unisys	Looking for longer distances at higher speed and a merger of SCSI and IPI

Vendors	
Adaptec	Remove SCSI limitations to expand the market
AMD	A sales opportunity to put Physical I/F into silicon
AMP	As a supplier of connectors interested in all aspects
AT&T	Offers a wide variety of transmission components and looking for places to sell high speed technology
BT&D	Provides components to support fiber applications
Fujitsu America	Expects to implement a system when the FC is defined
Hewlett Packard	FC is an application of fiber transmission
Integrated Photonics	IP sells fast fiber parts, and can contribute to developing and designing the FC
Potter & Brumfield	Has chips available to support fiber applications and serial to parallel conversion

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Components

John Lowell (BT&D) had brought material that covered transmitters, receivers and the integrated optics available in the market (a transmitter/receiver pair for FDDI is about $500 in quantity, dropping to $100 in 3-5 years).

In his overview, John introduced terms such as spectral wave length, cooler, modal spectrum, modal noise, sensitivity, dynamic range, and dispersion. A direct byproduct of this summary was a plea for a primer on the meaning of these terms. Glenn Moore (AMP) volunteered to have an applications engineer cover material which had been prepared to introduce fiber optic technology to electronics engineers.

John spoke about the most common source of laser failure not being the laser itself but the cooler, an integrated Peltier Thermoelectric Device which is designed to lower the laser temperature.

Del questioned John's MTBF claims of 1,000,000 hours at room temperature and 450,000 hours at 25 degrees Celsius, and asked if they were calculated or measured. The general feeling was that the MTBFs were unrealistically high. In a computer application, components operate at high temperatures and laser life depreciates relative to temperature.

Michael Pugh (Integrated Photonics) held himself in check remarkably well. The irritation Mike felt as he listened to some of John's statements was very obvious - he felt they were biased towards BT&D technology, and not necessarily representative of the industry.

Sensitivity was a new term which Mike explained. The higher sensitivity is, the longer the distance that can be supported. Unfortunately, that may mean the receiver cannot handle the overload condition that would exist at zero distance. Since telecommunication products have the goal of long distance, sensitivity is maximized.

In the data world, distance is so short it is more like a zero transmission length. The light does not dissipate enough, so the receiver must be capable of handling the full load. This will limit FC applications to using lower sensitivity receivers than normally used by communication companies.

John initiated a short discussion of integrated optical devices which could multiplex different transmission rates. One way is by simultaneous transfers in opposite directions and the other is with different optical wavelengths.

Reinhard Knerr (AT&T) listed all the reasons why simultaneous transfer on a single fiber is not realistic, even though it may be an appealing technique. The couplers at the end of a link cost 4 dB each and introduce directivity problems (the receiver picks up signals due to its close proximity to the companion transmitter), and reflection problems of the connector back into the receiver. Reinhard's description was sufficiently discouraging to kill all interest in simultaneous transmission.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Although multiplexing was interesting, suitable components are limited in availability, and very expensive. Estimates ranged as high as $3,000 each. Schelto Van Doorn (Potter & Brumfield) knows of at least one wavelength division multiplexer as low as $300. Another interesting but useless concept for the FC was the "ether loop" - a passive coupler that allows the mixing of television, phone and data communication on a single cable.

Schelto described the 128-pin 18-bit parallel-to-serial converter ECL chip that was designed for Siemen's medical group. There are 16 data bits plus two that are user-defined to indicate whether data is valid and if it is the high or low order half of a 32-bit transfer.

The link is a duplex channel which uses a scramble code embedded in the bit stream (which changes depending upon the content of the data) to produce a DC-balanced code over 10-bit intervals. The 18 bits of data are scrambled into 22 bits, for an overhead of 18%. The chips can be ganged together to operate synchronously if wider transfers are needed, sell for $550, and drop to $150 in volume. Clock recovery is external and there is no control - the link is permanently active and one of the two user bits is used as a Data Valid signal.

Mike Pugh presented Toplinc, which is currently being sold as an extender between HSC boxes. A module set costs $2000 in Q250, so the total for a link which can transfer at 50 MHz is approximately $4,000. The lowest achievable price point for the technology used is probably $1,000. Assumptions included in the price estimate are 50 µM or 62.5 µM multi-mode fiber with 8 nsec of skew on 12-fiber ribbon, and no more than 5 nsec of skew between each fiber. The length of a link at present is 1 KM.

Requirements

A potpourri of requirements was generated with no particular sequence or refinement, just a listing of what everybody thought should be relevant.

Cost Cheap (SCSI low end) Competitive w/copper Premium acceptable Speed 5 MBs 10 MBs 50 MBs 200 MBs Distance – Short (5M) – Long (2 KM) Error Rates Detected: 10^{-9} Undetected: 10^{-16}	Performance Low Overhead 80% Data Payload Low Latency (datagram) Minimum Station Delay Burst/Packet Size tbd Standby Power Sequenced Power Up to 32 Connections 200K Hours MTBF Connector Footprint < Copper Pigtail Ribbon Fiber	Environment 0-65° C 10-99% Humidity 0-5V but < 1A Shock Resistant Tight Cable Bend Applications Workstation Mainframe Fault Tolerant Redundant Paths Passive Bypass Transfer Integrity ECC or CRC

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Reinhard asked for background on the interfaces which would use the optical technology. A summary of the capabilities of SCSI, IPI and HSC was presented as a quid pro quo exchange for the optical information.

An attempt was made to put bounds on the requirements by defining the application areas, technology and cost. The market need is distance. In IPI and HSC it is for all electrical classes whereas for SCSI only the high end of shielded differential is seen as being able to afford fiber. Low end SCSI applications will never use fiber unless there is a dramatic reduction in cost. This might be achievable if consumer products adopt plastic fiber.

The combinations of speed and distance fell into broad categories of:

	< 10 MBs	> 10 MBs	> 100 MBs
< 100 M	✔		
< 500 M		✔	
< 2,000 M		✔	

Technology

SLEDs (Surface Light Emitting Diodes) are cheaper than laser components and there was a general consensus that ELEDs (Edge LEDs, sometimes referred to as failed lasers) are of little interest. Coherent mode operation is seen as impractical at this point in time.

The technologies were summarized in chart form.

Fiber	Mode	Transmitter	Wave Length nM	Freq/Distance nM	Length KM
Plastic	Multi	SLED	620	0.5	0.1
Glass	Multi	SLED Laser	800 1,300/1,550	500 1,000	5 10
Glass	Single	Laser	1,300/1,550	10,000	25

Although the plastic fiber can nominally operate at 100 M, the transfer rate falls abysmally so the practical limit is more like 10 M. Silica was also mentioned as an alternative material but there was little enthusiasm for it.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

The next part of the exercise was to put some kind of dollar estimate on the cost of a link in various technology combinations. A link was defined as what lies between the parallel interface on one circuit card over a 100 M length of cable to the equivalent component(s) on the other circuit board.

Speed MBs	**Distance** Meters	**Present Cost** $ in 3Q88	**Design-In Cost** $ in 3Q90	**1992 Production Volume**
1	< 500	?	< $ 25	1,000
12.5	2,000	≤ $ 2,000	< $ 600	200,000
50	1,000	≤ $ 4,000	< $ 1,000	50,000
50	> 10,000	≤ $ 6,000	< $ 2,000	25,000
125	> 10,000	≤ $ 10,000	$ 5 – 20,000	< 100

Perspective

The working group produced a successful exchange of information between the optical and interface technologists. A number of areas in which further information is needed were identified, such as practical experience (ICL MacroLAN), recent introductions (Hewlett Packard and National Advanced Systems), pending technology in other areas (IEEE 802.6 and SONET), and parallel fiber transfer components.

More exploratory work has to be done and the Functional Requirements need to be tied down before inviting proposals. One of the biggest problems may be too much interest. There could be no worse fate than to wind up as a battleground for those who did not achieve their goals in FDDI.

X3T9.3 IPI MEETING JUNE 20–21 1988

AMD hosted the meeting at the Embassy Suites in Milpitas. The atrium in the center of the hotel seemed awash with attendees, both visitors and locals.

Company	**Name**	**Company**	**Name**
3M	R. Herron	CDC Canada	R. Cummings
	J. Knudson	CDC MPI	T. Leland
AMP	C. Brill	Cinch Connectors	J. Guennewig
	R. Whiteman	Cipher Data	K. Geist
Ampex	L. Shih	Cray Research	J. Renwick
Amphenol	J. La Shier	Data General	T. De Osaran
AT&T Bell Labs	J. Morris		S. Solomon
Bull	P. Schreck	DEC	M. Halvorson
Burndy	H. Piorunneck		D. Shoemaker
CDC	W. Sanderson	Dilog	M. Nistanaki

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Company	Name	Company	Name
ENDL	D. Allan	Methode	R. Masterson
ETA	W. Hohn		D. Poplawski
Fujitsu America	R. Driscal	NEC Info Systems	J. Zilinski
	J. Luttrull	Northern Telecom	D. Wehrman
Fujitsu Compnts	R. Thornton	Scientific Comp Sys	D. Deel
Gould CSD	C. Pheterson	Stewart Connctrs	D. Hatch
Hirose Electric	M. Kosmala	StorageTek	D. Appleyard
Honda Connector	D. McFadden	Sun Micro	S. Nishtala
Honeywell Bull	D. Barney		S. Tsai
IBM	H. Brandt	T&B	H. Waltersdorf
	H. Truestedt	Tandem	J. Smith
ITT Cannon	W. Losch	Texas Instruments	A. Wetzel
JAE Electric	C. Salamon	Thinking Machines	E. Parish
Key Comp Labs	O. Gehlert	Ultra	N. Perdue
Los Alamos Lab	D. Tolmie	Unisys	R. Bergey
	R. Thomsen		A. Carroll
Madison Cable	R. Bellino	Viking Connctrs	K. Kwiat
Maxtor	D. Banerje		

Fiber Channel

Del Shoemaker indicated that there appeared to be opposition forming within X3T9.5 against the Fiber Channel on the basis that if such an interface was needed, SCSI and IPI should use FDDI.

Del made it clear that it was incumbent upon the committee to distinguish (and distance) itself clearly from FDDI.

This was done in re-drafting the SD3 to expand it to cover SCSI and HSC.

Duane Barney (Honeywell Bull) and Wayne Sanderson had generated a list of requirements they felt IPI users wanted to see met by the Fiber Channel. The only change made by the committee was to place a limit on addressability of 256 units, with up to 32 able to be physically attached.

General
Extend Host-Slave distances Support multi-Master operation Retain performance of copper IPI-1 Efficiently support Slave-Slave operation Provide smooth migration path from IPI-3 on copper

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Compatibility
Ascertain if IPI-2 support is needed Be usable by all device classes of IPI-3
Topology
Ring with a range of up to 256 addresses Distance supported to be optimized to 400 M Maximum distance between farthest units of 2 KM Redundancy supported by separate rings (optional) Provide attachments of up to 32 masters and/or slaves Continue operating when two adjacent units are not operational
Performance
Peak transfer rate of at least 10 MBs/fiber Desirable peak transfer rate of 25 MBs/fiber Effective transfer rate to approach peak transfer rate Use multiple fibers to multiply transfer rates (optional) Low power circuits to allow units to be left under power continuously

FIBER CHANNEL STUDY GROUP AUGUST 11–12 1988

Peter Dougherty (Unisys) had originally intended for the study group to meet onsite in Santa Clara, but arrangements fell through and the meeting wound up being held at the Sunnyvale Hilton.

Company	Name	Company	Name
AMD	J. Kubernic	Gazelle µCircuits	C. Popat
	P. Scott	Hewlett Packard	R. Haagens
AMP	C. Brill		D. Hanson
	J. Kevern		S. Joiner
	S. Yingst	IBM	P. Legakis
AT&T	M. Kao		R. Soderstrom
	R. Knerr		H. Truestedt
	J. Morris	ICL	T. Salthouse
CDC Canada	R. Cummings	Lasercom	G. Griffith
CDR Systems	M. Brooks		S. Sando
Ciprico	S. Reile	Lockheed	B. Peebles
	W. Winterstein	Los Alamos Labs	D. Tolmie
DEC	T. Martin	MMD	P. Ralston
Du Pont	M. Kaplit	NCR	J. Brown
ENDL	D. Allan	PCO	W. Philipson
Fujitsu America	R. Driscal	Protocol Engines	G. Chesson
	J. Luttrull		S. Cooper

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Company	Name	Company	Name
Protocol Engines	L. Green	Tandem	P. Sinykin
Signetics	N. Damouny		J. Smith
Sun µSystems	S. Joshi	Unisys	R. Bergey
Tandem	A. Pauker		P. Dougherty

Protocol Engine

Greg Chesson (Protocol Engine) spoke of designing silicon to handle Layers 3 and 4 of the OSI (Open Systems Interconnect) model by scaling a protocol processor so that it can handle transfer rates from 100 Mbs to 1 Gbs over multiple data paths. The objective is to do all the processing associated with a packet within the arrival time of that packet.

A 10 µsec budget has been allocated to handle FDDI, which handles a packet size of 128 bytes or larger in a continuous flow. It turns out that despite IPI (Intelligent Peripheral Interface) and HSC (High Speed Channel) being capable of 100 MBs transfer rates, the same budget is viable because both are parallel transfers.

It is possible to shrink the budget in a closed system to sub-microseconds by limiting the address range. Most of the budget is chewed up by general purpose address processing, so limiting the address range reduces overhead dramatically.

Protocol Engines is an incorporated joint venture in which Silicon Graphics is contributing the research and design effort. Revenue to support the staff comes from the $25,000 per annum contributions of TAB (Technical Advisory Board) members such as Apollo, Ardent, Artel, Ameritech, Brooktree, Martin Marietta, Unisys, Xerox, Interphase, and Mitre.

A primary objective beyond the immediate goals of developing silicon is to standardize XTP (eXpress Transfer Protocol) under the Navy's SAFENET program and within ANSI X3S3.

Existing communications protocols do not handle realtime requirements at all well, because in an activity where speed is of the essence, all the work has to be done in software. There are far too many connects (handshakes) needed, creating processing overhead and turnaround delays.

One advantage of XTP is its use of datagrams to avoid connects. Another one comes from moving the packet processing overhead to parallel data transfer instead of requiring it to be completed prior to initiating data transfer.

The major benefit of XTP comes from pipelining packet processing with data transfer. To do so requires the sequencing and coordination of independent state machines dedicated to a single task, such as the lookup match of the header. By the time the buffer is filled with the data, the overhead processing is completed. Committing to the "other" side of the data being routed is overlapped with the succeeding data packet arrival.

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

The five major elements pipelined are:

1. Header processing (address lookup)
2. Loading the context
3. Buffer filling and emptying
4. Protocol translation
5. Committing

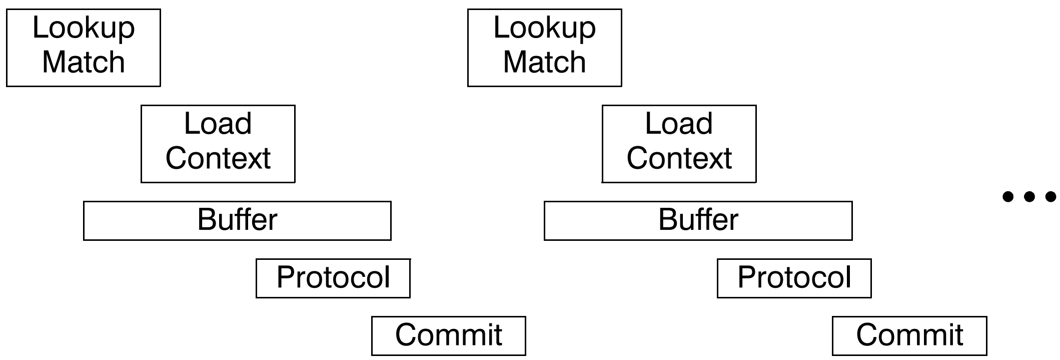

A growing problem, which will become even more acute in the future, is the linking of high speed networks in one location to others which are distant. Gateways will have to be constructed between LANs (Local Area Networks) such as FDDI, MANs (Metropolitan Area Networks) and WANs (Wide Area Networks). To succeed, compatibility with IEEE standards has to be emphasized.

XTP is a lightweight protocol, which goes beyond the strict needs of the message system to include system design considerations such as:

- Bulk Transfer
- Datagram support
- Reliable multicast
- Write control to avoid congestion
- Partial retransmission for long messages

The chip set design will be scaled to the technology available for silicon embodiment. It its ultimate form, it might be 4-chip architecture with a 32- bit microprocessor on each chip. In the interim, it is more likely to be a 3-chip implementation and a commercial microprocessor (multiplexed as needed between the three). The three chips will have high hardware commonality, and vary based on being programmed for a specific task.

Greg showed a sample of the C code needed to receive a packet. Using the RISC processor from MIPS as a case in point, the code takes fewer than 250 instructions, about 1,700 bytes. Improvements in XTP algorithms hold the promise of even more compact code.

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Jim Kubernic (AMD) asked how Protocol Engines expected to make money. Greg's response netted out to taking advantage of opportunities which are expected to arise in the infrastructure after XTP becomes publicly accepted as a standard e.g. assisting manufacturers in their implementations.

Roger Cummings (CDC Canada) recognized similarities between IPI and XTP and suggested a subset to marry IPI to networks. Greg admitted being ignorant of IPI but thought it might not be wise to be so close as to become involved in network management: stay pure and look after the closed system environment.

Bob Bergey (Unisys) was more interested in reducing the chip set by getting it down to a single part and Greg's response was associated with pin count. As the XTP has a 32-bit data path, 18 pins could be saved immediately by cutting back to 16-bits plus parity. Greg agreed to look at an IPI physical specifications before committing any further.

Missing the Boat

Greg had been invited to speak about the XTP protocol, but instead he made a report on program status, and then described the chip design and functional elements. It will be necessary to have him return to make a presentation on the protocol itself.

The problem arose because Greg had aimed his presentation at bit twiddling and chips. He has been reading the minutes of the HSC (High Speed Channel) working groups and was under the impression that this was the focus of the Fiber Channel also. It was pointed out that the Fiber Channel's objective is to come up with a physical interface that will protect the installed base of SCSI and IPI-3 software at the command level.

To clarify the difference the following simplification was used to explain how the three compare with what is needed:

		SCSI	IPI-3	FO Channel	HSC
Logical	Commands & Responses	Yes	Yes		
	MSG Protocol	Yes			
Physical	MSG Protocol	Yes			
	Acknowledge Protocol	Yes	Yes	Yes	
	Send Protocol	Yes	Yes	Yes	Yes
	Connectors, Cables, & Driver/Receivers	Yes	Yes	Yes	Yes

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

SCSI and IPI contain a complete set of functions. They vary in that the SCSI MSG protocol is a mixture of physical and logical functions which need to be layered more precisely. The HSC is not considered a channel as being simplex it is only half of a physical interface, and has no concept of acknowledging the receipt of either commands or data.

The purpose of the Fiber Channel is be a physical interface replacement for SCSI and IPI, which requires send and acknowledge closure, although it need not be in the same manner as is done in the current designs.

During his presentation Greg had made a statement that 75% of the throughput problem was not associated with the protocol, and that XTP was addressed at speeding up the remaining 25%. It was later discovered that the number of interpretations on what he meant by this remark was as large as the number of listeners. Everybody had a different understanding of what he meant, so a clarification is anticipated when he next attends.

Optical Primer

Jim Kevern (AMP) gave an introduction to fiber optics in which he alternated material between overheads and 35mm slides. As there was only one position in the room to position the projectors, Jim became rather accomplished at:

1. remove the overhead projector
2. move the table
3. put slide projector on table
4. focus slide projector
5. show slides
6. move table back
7. remove slide projector
8. put overhead projector on table
9. focus overhead projector

The circumstances could have been worse. If there had been only one female plug at the end of the extension cord he would have had to plug in for power between each motion. Fortunately for Jim, the hotel had supplied a cord with a multiple adapter plug.

Jim walked everyone through an explanation on the wavelengths of light and showed that the points chosen for light sources (800, 1,300 and 1,500 nM) are not accidental but based upon light properties. Jim's explanation on sources of refraction and indexing of fibers left all with a better appreciation of optical characteristics.

Red light travels faster than blue light and the waveguide is dependent on the wavelength of light. As light travels down the fiber, that amount which travels down the center will arrive at the opposite end before the remainder which bounces off the interior of the cable. To avoid this effect, the fiber is graded so passage down the center is slowed; refracted and center light travel together. Lasers do not suffer as much from dispersion as do LEDs.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Many of the terms used in optical fiber transmission are foreign to those who are familiar with copper. An example is Rayleigh scattering. Transmitted light always has some percentage of reflection back from the media to the transmitter. The action is referred to as backscatter, the effect is known as reflectance, and an Optical Time Domain Reflectometer (must be expensive with a name like that) is used to measure the loss.

Reflected energy in a source of signal loss in fiber transmission, and any interruption or discontinuity in the fiber exaggerates it. This includes connectors, splices, and even a tight bend in the cable.

There is an order of magnitude difference in the effective impedance between copper and fiber. Fiber starts with low power. It gets worse as attempts to tap it produces large losses. Multidrop configurations are not feasible.

There was much interest in the potential of plastics. Jim is not optimistic about its wide use. Del Hanson (Hewlett Packard) volunteered the information that DAT (Digital Audio Tape) decks use plastic fiber for noise isolation.

Lasers

Ron Soderstrom (IBM) had also prepared an introduction on fiber optics, so he brushed past the duplicated content quickly. He summarized the advantages of fiber as being longer transmission distances at higher data rates than copper can achieve.

Optical fiber cable has been at a point of technology development where few improvements can be expected but the fiber transmission components needed to support the cable have changed dramatically in the past few years, and are continuing to improve.

Driver Electronics	Receiver Amplifier
Optical Source	Photodetector
Lasers	APD Photodiode
LEDs	PIN Photodiode

Copper and optical transmission systems vary because of the need to convert electrons to photons and back again.

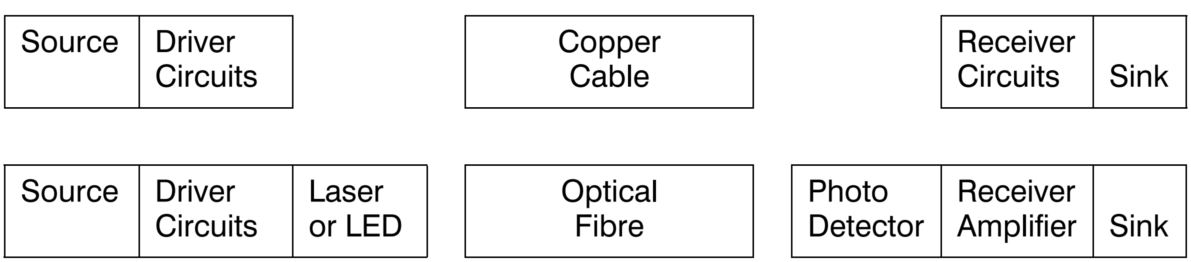

Optical fiber has to be transparent at the selected wavelength, thus there are only two choices of material, glass and PMMA (PolyMethyl MethAcrylate) plastic.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Plastic is good for relatively short distances. It has a large core of 1,400 µM compared with 50 µM for glass, with the advantage of a very high acceptance angle which must be traded off against low bandwidth. It can be a very economical transmission system using low cost LEDs.

Glass offers the highest transparency, and is based on silica. There is the potential that fluoride glass will complement silica at some future time, to offer higher performance because of its greater wavelength.

SLEDs (Surface LEDs), ELEDs (Edge Emitting LEDs) and lasers are usable light sources. The problem with an ELED is stopping it lasing. There are major differences in property between lasers and LEDs:

Lasers may require a high threshold of initial power to begin driving, but modulation requires only a fraction. LEds have no threshold, so modulation requires the full voltage swing. Lasers also significantly reduce the effect of dispersion, with a spectrum of only 2 µM vs >100 µM for LEDs.

		Laser	LED
Optical Power			
Output Power	(mW)	1.0 – 40.0	1.0 – 10.0
Fiber Power	(mW)	0.5 – 5.0	.1 – 0.3
Coupling Efficiency	(%)	10.0 – 70.0	1.0 – 5.0
Output Spot Diameter	(µM)	1.0 – 3.0	30.0 – 150.0
Surface or Edge Emitters		Edge	Both
Response Times			
Rise/Fall Times	(nsec)	0.02 – 0.5	2.0 – 5.0
Transfer Rate	(MBs)	0.01 – 20.0	0.001 – 0.2
Drive Requirements			
Forward Current	(mA)	10.0 – 200.0	50.0 – 300.0
Threshold Current	(mA)	5.0 – 170.0	n/a
Modulation Current	(mA)	2.0 – 30.0	50.0 – 300.0
Voltage Drop	(V)	1.0 – 2.0	1.5 – 2.5
Spectral Properties			
Wavelength (GaAIAS)	(µM)	0.78 – 0.86	0.8 – 0.9
Wavelength (InGaAsP)	(µM)	1.3 and 1.55	1.3
Spectral Width	(nM)	0.01 – 2.0	70.0 – 200.0
Coherent		Yes	No
Double-Hetrojunction		Yes	Yes

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The two photodiode detector types are APD (P-N junction with Avalanche Multiplication) and PIN (P-N junction with Intrinsic region). Both are available for all fiber windows (800, 1,300 and 1,500 nM).

PIN	APD
Low Bias Voltage (5-90V) Low Cost High Volume (Mature) High Reliability	High Bias Voltage (150-300V) High Cost Low Volume High Receiver Sensitivity

There are a limited number of ways in which network configurations can be configured: ring, star, switch and mesh structures.

Ring	Star
Token on insertion Extendable to many nodes Peer-Peer communications Active or passive T-couplers Simultaneous messages possible Optical bypass switch disconnect	Passive or active Multiport connections Relatively inexpensive Needs collision arbitration One message possible at a time Avoid collisions by using tokens
Switch	**Mesh**
Multiple messages/high throughput Connect any node to any other node Expensive switch optics/electronics	Difficult to reconfigure Complete interconnection Point-Point link between each node

Data communication requirements differ greatly from telecommunications need for long haul, inner city loop and fiber to the home. Fiber to the home is being installed on an experimental basis in over a dozen field test sites. This means that the local distribution loop will eventually bring single mode fiber to every home and every business in the country.

Telecommunications	Data Communications
Single mode fiber Laser sources: 1,300 and 1,500 nM Long distances of 10–100 kM	Single or multimode fiber Laser sources: 800, 1,300 and 1,500 nM Short distances of up to 4 kM

A natural byproduct of this explosion in fiber use is that 1,300 nM lasers will become high volume, low cost components. Data communications can build on the leverage such a high volume provides and decrease system cost. It is feasible for both telecommunications and data communications to use the same fiber lines.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

At present, initial installations are of POTS (Plain Old Telephone Service) with future potential for offering video, pay TV, mail, and teleconferencing market opportunities.

With this kind of background, it is easy to justify interest in single mode fiber and lasers for low cost, universal application. As a source of light, lasers are far superior to LEDs.

- Lasers have a higher power budget
  - Star couplers and splitters are feasible
  - Can have tailgate and back-panel connectors
  - More connectors and longer distances are practical
- Lasers promise lower system cost
  - CD lasers are a high volume consumer product component
  - CD package is a simple construction (automated assembly)
  - LED packages are complex electronics and optics assemblies
- Lasers have 100 times less Transmitter/Receiver Crosstalk
  - Laser drive current is 10 times smaller than LEDs
  - Laser receiver sensitivity is 10 times less sensitive than LEDs
- Lasers have superior detection and fault isolation capability
  - Both LEDs and lasers have a fault signal at the receiver
  - Lasers have a fault signal at the transmitter (back facet monitor)
- Lasers can be focused to a smaller spot diameter
  - High coupling efficiency (>50%) into any size multimode fiber
  - Good coupling efficiency (>25%) into any size single mode fiber
  - Poor LED coupling efficiency (>5%) into large core multimode fiber
  - Laser spot diameter is only 1–3 µM compared with 50–150 mM for LEDs
- Lasers permit faster transfer rates
  - LED rise/fall times are 200–500 nsec (up to 200 Mbs)
  - CD ROM laser rise/fall times are 200–500 psec (up to 2 Gbs)
  - Communication laser rise/fall times are only 20–70 psec (up to 20 Gbs)

Ron stressed the advantage in a system application of being able to monitor light reflected from the back facet with a photodiode, as it both simplifies and improves the diagnostics capability.

CD ROM players had to use lasers because such a small spot size is needed that LEDs could not cut it. Volume on the CD ROM laser is several million/month. The cost has fallen to less than $4 apiece. They come in a compact, hermetically sealed 5.6 mm diameter package, are compatible with a single +5V supply, and operate at 780 nM. A change in the wavelength of 780 nM is inconceivable because the depth of the pits is a quarter of the wavelength. The fast 200-500 psec rise/fall times hold the promise for transfer rates of up around 1 Gbs.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

In answer to his own rhetorical question; '*If lasers are so good how come nobody has used them?*', Ron answered that the biggest problem in the past has been reliability concerns. This is now disappearing as volume production resulted in manufacturing processes that simultaneously had the effect of improving both yields and device reliability.

To date, IBM has accumulated more than 10 million hours of testing on over 1,000 lasers and is continuing to accumulate data at a rate of almost 1 1/2 million hours/month. The slope of reliability improvement has been virtually a straight line since the initial evaluations made by Bell Labs in 1976 to IBM's latest results. The line looks as though it may continue falling until the early 1990s, by which time, the present laser MTBF of 1,000,000 should be around 5,000,000 hours.

The primordial cause of laser failures is heat. Laser reliability decays rapidly at elevated temperatures, and is a stark contrast to LEDs which are insensitive to this effect.

Lasers can be tested to find out how likely they are to fail. The value of T0 (the amount of current needed to drive a laser) is already tested by CD ROM manufacturers at two temperature points. If the difference in current required is close, this represents a high T0 and implies an extended life. Empirical observation has proved that laser life will be relatively short if there is a large gap in current required.

Other parameters have been tested at IBM. Periodic checks on characteristics such as spectral response and spectral width have shown no detectable change over life. The failure rate at a case temperature of 50 degrees Centigrade shows a median life of about 1,000,000 hours.

The 1,300 nM lasers are doing about 10 times better than 1,500 nM lasers. All testing is done under continuous wave i.e. constant emitting. Modulation does not seem to affect the life of the laser at all, and thermal cycling appears to have no effect on performance.

Photographs of the test equipment showed racks of 128 lasers which can be closely monitored, and others of 1,024 that can perform fewer tests on each laser but can cycle the temperatures to predict end of life.

A CD ROM laser used with single mode fiber actually operates as double mode in the popular single mode fibers available today because of the high modal dispersion. Since the two modes would interfere at the receiver, one has to be eliminated before transmission, but this is a minor drawback to the advantages that can be obtained.

It is impossible to compare a Tx/Rx (transmitter/receiver) package for FDDI with a CD ROM laser as the FDDI package contains the mechanical housing, aligns the laser, couples the lens and fiber connector together and includes clock recovery and some other basic circuits. All of this must be external to the CD laser, but it does not require much real estate on a board.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Note: IBM has selectively shown a small card with transmitter on one side and receiver on the other. The PCB (Printed Circuit Board) ground plane acts as the shield to prevent crosstalk and there is only a few components to fit everything together. The cost of an FDDI Tx/Tx is expected to be $150 in 2-3 years and at $4 for a CD ROM laser there is $146 left over.

This can go a long way to covering the cost of cobbling together the extra bits-n-pieces needed for a 3:1 performance improvement over FDDI.

IBM's card looks to have far less than $50 worth of parts on it, but it is not a trivial task to get it to all working properly. Ron admitted to having required several iterations of silicon to get clock recovery and the other features needed operational.

Since the basic ground rule was low cost, no modifications to the CD laser package was allowed. A standard fiber output connector was used instead of a pig tail. Everything was achieved with off-the-shelf components, even the grin lens used to couple the laser output into the fiber.

The low cost silicon PIN photodetector used has 500 MHz of bandwidth with a large active area of 0.25 sq mm, with no coupling optics required. A transimpedance amplifier is used to eliminate the need for equalization and there is a wide dynamic range (>20dB). The available signal power is 20 dB, almost twice that of FDDI (-3 dB launch power and -23 dB receiver sensitivity).

The BER (Bit Error Rate) is better than 10^{-12} at the minimum receiver sensitivity. Although 500 Mbs is feasible, the problem in achieving this transfer rate is support silicon. The present link is limited to 300 Mbs on 2 KM of fiber, but appears capable of 400 Mbs. Tests made by researchers outside of IBM have achieved higher rates, using "*super silicon*" to reach 1 Gbs.

The push for low cost has Ron looking to the recently introduced SC plastic connector which uses a push on/press pull latching mechanism instead of the older machined metal FC connector that screws on. An added advantage is that connectors can be more closely packed together on the bulkhead.

SCSI Fiber Extender

Marc Brooks (CDR Systems) described Paralan, a serial extender for SCSI that has been implemented in coax and is going to be offered as fiber. Two black boxes are needed, one at each SCSI site to be connected. The SCSI protocol is converted to a proprietary bit serial protocol between the black boxes.

Marc refused to give any technical details. It was deduced that by transmitting continuously with a protocol that does not need to wait for the turnaround in a connection, CDR Systems has halved the turnaround delays of the SCSI protocol. It was still unclear as to how it all works but as Mark was unwilling to explain further, the subject was dropped.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The coaxial black boxes are $395 each and the fiber optic ones are $790. The maximum distance is 300 M with either RG62 coax or 62.5 μM fiber. A single fiber is used. Full duplex operation uses split LED transmitters/receivers. The first test site is a Los Angeles bank.

Initial market response to product announcement has been very enthusiastic, with most installation requests being around the 300 M mark. Patent applications have been made on many of the Paralan features and there is no intent to propose it for standardization at this time, but to keep it as a proprietary advantage.

DQDB

Del Hanson described the principles behind the IEEE 802.6 DQDB (Distributed Queueing Dual Bus) project. Relations between X3T9.5 and 802.6 have been acrimonious for several years, with disputes over charter.

Matters escalated last year over the previous 802.6 project which "borrowed" the FDDI MAC (Fiber Distributed Data Interface Medium Access Control) and made several incompatible modifications to it.

DQDB originated as QPSX (Queued Packet and Synchronous Switch) in Australia, and the name would have been retained but the project was incorporated as a company under the same name. The roots of QPSX can actually be traced to an experimental project at AT&T Bell Laboratories which became a joint effort with the communications authority in Australia. Its a case of having gone full circle, as it is now back with AT&T support at IEEE.

The claimed advantage of DQDB over ring-based architectures such as FDDI and FDDI-2 is that node failures do not require system reconfiguration because the dual buses are unidirectional. DQDB does not have point-point clocking like a ring but a master clock, usually located in the telephone network. NOTE: This is claimed to be superior to FDDI but since both techniques have their pros and cons it should really be noted as a difference.

Typically, DQDB will be configured like a ring with both the master and the slave frame generator in the same cabinet (packets received on the master have to be re-generated on the slave). Why this is claimed to be more robust than FDDI confused the group, as the frame generator is a single point of failure that can bring down the system. Even if master and slave generators are in separate cabinets, the loss of either would kill the system.

Nodes tap into the bus, check the header of passing packets, and only if they match, pull off the data. *Note:* This too is claimed to be superior to FDDI, but as the physical layer of a node is responsible for repeating the slot contents and a physical failures requires reconfiguration of the buses, this too is more of a difference than a feature.

The bus works with fixed time slots of 125 μsec. A packet is made up by concatenating as many slots as are needed; the number being a function of the transfer rate of the PMD (Physical Medium Dependent) sublayer adopted by an implementation. The overhead to concatenate slots is about 7%.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Note: Most of the claimed superiorities of DQDB over FDDI can be attributed more to religious preference than technical merit, as one engineer's feature is another engineer's deficiency. DQDB does differ in one important respect, each node keeps track of network loading.

Headers have Busy and Request bits. When a node wants to transfer data, it sets the Request bit in as many slots as it needs on one bus, and when it transmits data, sets Busy in the slots that it is using on the other bus. A node which signals on one bus, transfers on the other, so it is best to view a node as being either left handed or right handed.

Each node maintains a Request Counter. The counter is incremented for each Request slot that passes and decremented when an empty slot passes by on the other bus. Decrementing the count is an implicit recognition that the slot "belongs" to a node further downstream. When the Counter is at zero, the node captures the first empty slots that come by, fills them with data, and marks them as Busy. The node strips its own slots when they return.

This technique provides faster access than a token ring, because access can be immediate if the Request Counter is zero vs an average latency of half the rotation time on a ring.

Parallel Connectors

Kaushik Akkapeddi (AT&T) described Bell Telephone's parallel fiber connectors which was designed for direct board attachment (as would be done in a VME chassis). Two models terminate 12 or 18 fibers, which may be ribbon or discrete.

The connector is an impressive exercise in technology. Parallel fibers are aligned in an etched silicon V-groove carrier which is mated with another to form a silicon sandwich. The face is polished for contact and the sides for position control. The corners are beveled to reduce the risk of chipping. Four springs are used to align the sandwich into the connector housing.

Although ceramic is the preferred material for single fiber ferrules, it is not practical as the sandwich material. It is impossible to achieve the same degree of accuracy with machining as it is with etching. The silicon etches along its crystal boundaries, so the tolerances are extremely tight. Testing has been done to 1,000 insertions with no problems encountered.

Ribbon cable looks fragile and appears only suitable for use internal to the cabinet but it is available sheathed for building installation. The ribbon fibers use loose tube construction so are free floating inside the cable and do not suffer from cracking with bending.

Price seems reasonable at approximately $1.20/foot for 12 unsheathed fibers. The 62.5 µM multimode fiber has less than 3 nsec/KM of skew, and it is at least 10 times better for single mode fiber.

The average loss over all of the fibers in multimode is 0.53 dB and it is 0.7 dB in single mode.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The cost of a mated pair of fiber connections is $25, so a 12-fiber termination is $300 per cable. This cost is expected to drop, as it is such an early point in the production cycle; only about 3,000 connectors are in field use today.

Improvements not yet made to production connectors include beveling the face of the silicon so that light will be dispersed sideways into the cladding, instead of causing possible cross-interference over the flat surface used at present. To ensure proper mating the receptacles will be keyed.

Although the connector costs may seem a little disconcerting at this time, it is unlikely to be prohibitive in terms of the total system cost.

MacroLAN

England's International Computers has gained more experience on fiber connections in the last four years than any other company. Over 20,000 installed nodes represent more than the sum of all others installed in the world. Tony Salthouse (ICL) described the project and the technology used.

MacroLAN was not introduced as an alternative to the channel architecture of previous systems but as a complete replacement for linking peripherals and processors together. All parts of the system were impacted by the decision.

Layering is seen by many as an obstacle to comprehension and implementation, but ICL found it to be a way of thinking that was absolutely essential to getting MacroLAN operational. A word of warning is not to underestimate the size of the integration task: it took tremendous resources (far more than were anticipated), to connect everything together.

Configuration management consumed development time. Seemingly simple decisions such as which CPU takes the dumps are not simple in a fully connected environment. Combining the peripheral and communication support in the same software drivers required a great effort.

The initial concepts developed in 1980 resulted in a re-invention of token ring. A visit from Soderblom (who claims to hold exclusive patents on the concepts of the token ring) changed all this to become what is referred to as a token bus.

Distributing the repeaters required for optical connections at the stations would mean having to provide expensive optical bypass switches to provide for power outages. It was decided to use active switch boxes (known as MAC-less concentrators in FDDI) which have redundant power supplies instead.

A station is not active unless it is transmitting, so inclusion or exclusion from the system is a decision made by the port switch. MacroLAN operates on the same principles as a token ring for control: the right to transfer data is provided by a token.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

MacroLAN becomes a bus during data transfer: information is broadcast on all lines by the port switch. It is a hybrid, employing a ring protocol for control and acting as a bus for data transfer.

No one station is a master and a datagram service is used for I/O. It came as somewhat of a surprise to learn that MacroLAN basically complies with the ISO OSI (Open System Interconnect) model. Data packets may be up to 512 bytes with a 5-byte header and 4-byte trailer. The Transport protocol is compliant with ECMA 72 Class 4; a forerunner of the OSI transport model in which every packet is acknowledged.

The transmission system is based on PIN diode LEDs at 1,300 nM. Graded index silica fiber is used in the "old fashioned" 100/140 (ID/OD) µM diameter. The large fiber size was felt to be essential for improved launch power.

Terminations are done both at the factory and in the field using a "glue and cleave" technology. After epoxy has set around the fiber a tool is used to cleave the end and leave a clean, optical interface. The repetition rate is 100 Mbaud and the Manchester bifrequency code achieves 50 Mbs. Bit stuffing is used to prevent a stream of too many 1s or 0s.

Customers have benefited greatly from the development of MacroLAN. Those who are looking for resilience can split a single system over widely dispersed sites e.g. on either side of a road. The user, for the first time, has been able to put peripherals wherever he wanted to rather than in close proximity to the central processor.

For those to whom security is very important, fiber meets Tempest requirements and eliminates EMI/RFI (ElectroMagnetic Interference/Radio Frequency Interference). From ICL's point of view, the biggest advantage has been that there are no rogue sites which have grounding problems requiring intensive CE (Customer Engineering) effort to locate. Other benefits still continue to appear, although the most obvious one may be best - it is physically easy to plug in.

When looking at moving to a fiber environment, good diagnostic aids are an absolute must to help isolate problems in the configuration. It may be difficult, but it is absolutely necessary to develop a culture in the customer environment of cleaning the connectors and treating them with more respect than the old copper types.

There are drawbacks to MacroLAN - it is too flexible and requires that a closed system implement something like LAN station management. Even after four years minor problems still keep popping up which need to be addressed and corrected.

The port switch permits up to 6 stations to be attached. It is a small wall box mounted close to a power outlet. The switch is a star distributor with a clock generator. Up to three switches can be concatenated to achieve a total distance of 1.5 KM. The maximum configuration ICL supports consists of 64 controllers and 16 CPUs.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special provision has been made for processor-processor communication. Rapid updates may be made for Test and Set type functions on the smaller configurations. Although MacroLAN is relatively fast, it has proven inadequate for the larger processor complexes which use 4 fibers for processor-processor communication using a function-specific unique protocol.

One consideration of the optical plant is that the volume of fiber used is relatively small, so its cost is not a big deal. The real need is to keep the cost of termination low. The connectors are sole source, unique, and the supplier is Thomas & Betts. The ruggedness of the cable has been improved by using three Kevlon strands in the sheath instead of two, so that there is better resistance to crushing.

Although a lot of terminations had been anticipated in the field, it turned out they have not been needed. The field installable connector has proven to be a requirement, despite the low frequency of connector damage.

The protocol is very fair. There is no contention, and as MacroLAN operates at very high efficiency no problems have been encountered in running five controllers with an average of six 3 MBs disks on each. The arithmetic sum of the disk transfer rate is 90 MBs, but by the time seeks and rotational latency are accounted for the sustained transfer rate is comfortably handled by the 50 Mbs MacroLAN.

The port switch is quite small at 15" × 9" × 3" and mounts easily to a wall. Cable abutment and passive connectors are not allowed because of the signal losses involved. ICL requires that if the user needs to extend cables that it be by a port switch so that there is an active repeater.

HP-FL

Randy Haagens (Hewlett Packard) described the different approach which was taken in developing FL (Fiber Optic Link). It was felt essential to retain an architecture as close as possible to existing channel architectures in order to reduce impact on the rest of the system.

A direct equivalent to the FL is embodied in the P-bus, a copper version for daisychaining within a cabinet. The first goals set for the project were not unlike those of ICL, but it was recognized that the size of the effort and its impact would be so large that the effort was scaled back.

The prime need was recognized as being the ability to locate disk subsystems farther away from the processor, so this redirected the project. The system is very conventional, limited to one processor and its attached peripherals.

A storage cluster contains 8 disk drives, each of which has an embedded controller with both FL and P-Bus connections.

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Note: The only justification for including the cost of both in every drive is to make it easy to interconnect the disks to multiple processors for loosely coupled multiprocessing.

The technology chosen is remarkably similar to ICL, using 110/140 µM fiber but with 820 nM LEDs (not 1,300 like ICL). All of the optical components are sourced internally to HP.

The original design goal of 100 M was exceeded - it turned out to be a 500 M link. A proprietary 5B/6B code is used. A special chip (not unlike AMD's TAXI for FDDI) was designed for parallel-serial conversion.

The FL can sustain a net transfer rate of 40 Mbs, plus flow control plus error handling. A packet consists of up to 16 data bytes framed by a 1-byte header and a 2-byte CRC plus a 1-byte end of frame delimiter.

Frame errors are recovered using a "*go-back-in*" algorithm i.e. recognize the error in transmission and begin transmission again from the failed point. The small packet size was selected because of memory restrictions in the flow control chip (multiple frame buffers are required in order to sustain the transfer rate).

Switches

During discussions on how to configure a system, Randy recommended that the focus should be centered around the physical plant rather than the logical configuration. The requirements approved at the August plenary included the criteria that the configuration be a ring.

Randy questioned the wisdom of this, and pointed out the benefits of being able to configure a physical plant using switches to simplify the management problems. He suggested worrying about the logical characteristics after the physical plant configuration has been decided upon.

This endorsed the approach taken by ICL, using what Tony referred to as the active switch. Tony objected to using intelligent as a term to describe the MacroLAN switch as it does not perform any functions of routing, buffering or multiplexing.

What was clear from the discussions is that there is a lack of consistency amongst attendees on interpreting what switching does. Randy agreed to put together a glossary of switching terms and definitions that could be used to clarify the syntax of discussions, so everyone can understand everyone else.

Del Hanson volunteered AT&T to make a presentation on fast packet switching, which is recognized as a leader in this area. Jim Morris agreed to do what he could to locate somebody to make a presentation.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Protocol

Both SCSI and IPI have what are described as connection-oriented protocols, because of the amount of handshaking needed. Except in data transfer, both stall during the protocol sequences until actions are acknowledged by the other end of the cable.

This may be acceptable at the short distances of copper interconnects but is quite impractical at the extended distances of optical fiber links. It will be necessary to use a protocol that does not require acknowledgment from the far end before the next successive step in a sequence is executed.

Under the FL protocol, a command transferred to the controller is assumed to be received, and it is acknowledged by command completion. If there is no response within a reasonable period, the CPU times out, assumes the command was not received and sends it to be re-executed.

Tim Martin (DEC) questioned whether under IPI-3 it is possible for the same command to be re-issued and have it executed without a change in result - an innocent question that produced a storm of confusion on what is in the IPI-3 protocol vs what is in the IPI-1 protocol, all mixed up with what is needed in the protocol of the fiber channel.

It took some time to clarify that the responsibility for ensuring that the media contents do not differ as a result of additional command transfers, is a responsibility of software initiating the commands, and not IPI-3. If a command sequence to read, followed by an update write does not reserve the disk, it is not the responsibility of IPI-3 to detect the data corruption which can occur if an intervening write command is received on another port.

FIBER CHANNEL STUDY GROUP SEPTEMBER 26–27 1988

The meeting was hosted by Hewlett Packard on Trimble Road in San Jose.

Company	Name	Company	Name
AMD	N. Mammen	Fujitsu America	B. Driscal
AMP	C. Brill		J. Luttrull
Ardent	D. Shaffer	Hewlett Packard	R. Haagens
AT&T	M. Kao		D. Hanson
	J. E Morris		S. Joiner
	M. Pashan	IBM	E. Johns
Avantek	G. LaBelle		H. L Truestedt
Canstar	K. Lue Shing		C. Wong
	K. Malavalli	Lasercom	S. Sando
CDC	W. Sanderson	National Semi	P. Sweazey
CDC Canada	R. Cummings	Optivision	A. Dias
ENDL	D. Allan		M. Hauke

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Company	Name	Company	Name
Optivision	A. Jain E. Onstead	Tandem Unisys	J. Smith G. Carter P. Dougherty D. Mortensen
Tandem	A. Pauker D. Pham P. Sinykin		

Packet Switching

Mark Pashan (AT&T Bell Labs) reviewed the principles of packet switching. The first implementations used computer memory, but this limited volume. The next evolutionary step was time division, the use of a single shared medium such as a ring or a bus. Space division is where its at today, with switch arrays, also known as intelligent crosspoints, in which packets are self routing and there is maximum parallelism.

The latter has brought the level of performance up from switching packets at a rate of tens of thousands/second to millions/second and cut the delay from 100-150 msec down to less than 1 msec.

AT&T has been active in the implementation of packet switches throughout all the stages and has experimented with several design concepts. For a video conferencing system, a Richards Network was used which has two stages; there may be several input terminals, but the number of outputs on each is a function of the number of nodes at the second stage. The whole design was based on circuit switching, and had the unique characteristic of being able to broadcast. Although a technical success, the crosspoints were very expensive and video conferencing never made it for economic reasons.

As one way of reducing system cost, the 'knockout' principle can be used. As a packet enters a switch it is either accepted or 'knocked out'. At some point, there is inadequate buffering to hold the packets not being transmitted; these are lost and need to be retransmitted. Each input is put on a separate broadcast bus and each output has access to all the input lines. The advantage is that only small buffers are needed.

The big breakthrough in performance came from the Banyan Network which does automatic address translation, a huge compute problem in large systems. Each packet is given an internal header which routes it through the fabric of the switch in a predetermined manner. With only one way for any packet to get through a switch, congestion anywhere on the route will cause problems. One way to reduce congestion is to have a packet buffer in each node.

The design is replicated for each node in a Banyan Network, and iterating on each bit of the internal header address requires identical logic. Priority routing is possible, as in the mixed voice and data switch built that ranked Signal Control packets the highest, then voice, then data packets.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The performance of a Banyan network can be improved further by sorting the input packets to reduce congestion. Collisions which cause congestion can be identified. The overflow packets are trapped and kicked aside into a buffer to be re-submitted to the input stream and try again, this time with a high priority attached.

In 1986 AT&T conducted a wideband packet technology field experiment in the Bay area. A mixture of media was used; coax, radio, T-carrier and fiber. An ADPCM (Adaptive Differential Pulse Code Modulation) technique was used which integrated the reduction in data (as packet volume built up) over all conversations. A typical DS1 link capable of 24 voice channels was able to carry 120 conversations before hiss and pop became annoying. About 2/3 of speech is silence and the remaining improvement came from the ADPCM data reduction.

New work is being conducted with optical switching technologies. These have all sorts of potential for switching packets even faster, but a description of them sounds like science fiction. Multichannel Multihop brings visions to mind of bunnies in a field but it refers to wavelength multiplexing; hopping from one wavelength to another inside the fiber.

Types of Switching

Randy Haagens (Hewlett Packard) complemented Mark's material with a summary of the different types of switching.

- Circuit Switching: Exclusive dedication of a fixed bandwidth between two users for the duration of a call
- Fast Circuit Switching: Switches can perform call setup and signaling so quickly that it becomes efficient to switch brief bursts of data
- Message Switching: Complete message is stored at each switch, then it is forwarded according to routing information in the header
- Packet Switching: Messages divided into smaller segments of fixed maximum length
- Fast Packet Switching: Like packet switching, but a simplified physical interface and smaller packets
- Burst Switching: Like message switching, but burst flows through switch without buffering

The fixed, dedicated bandwidth of a circuit switch is inefficient because of the slow overhead associated with call setup. No practical implementation of a fast circuit switch has yet been made, but it is an ideal that excites telephone companies because it makes significantly more efficient utilization of the bandwidth.

Message switching is obsolete, because the cost and overhead of holding the complete message in memory is both expensive and inefficient. It adds delays in repeating the packet because nothing is repeated until the end of the message has arrived. Packet switching took over because it is more efficient to break up the message into smaller units and repeat data in granules.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

A peripheral channel is a burst switch which has simplified routing (limited number of direct addresses). Control overhead is low and applicable to burst size. Queueing to gain access to the channel is the source of delays, which is variable, and typically is short.

Objectives

Randy also provided definitions of what he felt the objectives of the Fiber Channel are:

- To design a subnetwork capable of connecting multiple SCSI and IPI to multiple processors
- To define a subnetwork link suitable for direct attachment of devices adhering to SCSI or IPI higher-level protocols

The choices are a ring or a bus, which have similar access methods such as insertion, token, or slotted vs the point-point star configuration which has centralized control. The latter is probably the closest in philosophy to the existing I/O structure.

One possible scenario described to merge the installed base into a FC (Fiber Channel) environment was to have SCSI and IPI bridge products which would convert parallel copper on one side to fiber on the other. The two would be incompatible in several respects e.g. the protocol for the fiber channel is going to be datagram instead of connection oriented.

Obviously, this is less desirable than a cohesive implementation such as the Fiber Link recently introduced by Hewlett Packard that has a 50-pin parallel copper equivalent internal to the cabinet named the P-bus. Wayne Sanderson (CDC) questioned the merits of developing new parallel variations on SCSI or IPI. He suggested instead, that a more effective plan would be to develop a copper version of the FC which would use less wires and smaller connectors than a parallel version.

The subsetting of the FC could be:

Serialized Fiber Channel	Optical to Electrical Conversion	**Serialized Copper Channel**	Serial to Parallel Conversion	**Parallel Copper Channel**

The parallel copper channel should represent a relatively small step for the existing SCSI and IPI, so that the installed base can easily identify with the new physical interface. There are several features desired in SCSI-3 which will affect the existing SCSI physical interface, and if new silicon is needed then it would be desirable that it be a subset of the FC. Systems integrators could use the FC for external connections and copper for those internal to the cabinet.

Wayne accepted the merits of this approach for SCSI but does not believe the same argument applies to IPI. The only advantage to a parallel version would be the datagram service, and this would not be a big enough step forward to justify its implementation.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.
A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.
This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

As a rough rule of thumb, it was assumed the signal count for a parallel implementation would be about 25:1 over fiber i.e. a 2-fiber FC would be equivalent to a 50-pin parallel implementation.

Paul Sweazey (National Semiconductor) felt that chips to support serialized copper would be cost-effective, in large enough volumes. Paul needed convincing that a fiber connection could not be cost competitive with copper even if the silicon costs were comparable.

Once Paul was convinced that fiber would never get to be cost-competitive with copper, Del Hanson (Hewlett Packard) questioned whether serialized copper is feasible. At the high repetition rates, EMI/RFI considerations may make a copper version impractical. This did not faze Wayne, who would look to coax or copper sheathed flat ribbon cable as a solution.

At SCSI meetings, Greg Floryance (IBM) was master of the innocent question which generated a storm of controversy. For FC, Armando Pauker (Tandem) assumed the role when, on the subject of granularity, he asked if it was acceptable to achieve a 200 MBs system with ten 20 MBs implementations. He was given a mixed chorus of Yes and No from the group.

Wayne responded on the purpose of a standard being interoperability and that if any variety of combinations was allowed then this objective would suffer. Randy felt some multiplicity should be allowed as increments in performance can be gained within the same technology but suggested that the limit be set to a reasonable number such as four. Paul joined in, as the best economics come from changing technology, instead of duplicating what exists.

Against these proponents of limiting the number of channels that could be repeated was the argument that such a decision more properly belongs to the systems integrator. In the typical situation, if the system need is for an overall increase in system bandwidth, then changing technology is the best choice.

However, if all the attached peripherals are happily within bandwidth on 10 MBs channels but there is one such as a SSD (Solid State Disk) which is ten times faster, then a cheaper system solution may be to add 9 cables to the SSD. Spares, support, training, development time and costs are all factors which have to be weighed.

Randy stepped back from the fray, waxed philosophical on the self limiting nature of technology, and drew the debate to a close. The benefits that are derived from stepping up to a faster channel are strong enough that systems integrators cannot afford the duplication except in rare instances. If the increment of improvement is 4-5 (as from 10 to 50 MBs or 50 to 200 MBs then the break point for duplication vs changing technology is probably the apex. Any more parallelism than 4-5 is uneconomic unless it is a special case.

Requirements

Most of the meeting time was spent on building a reference model, or list of requirements, for the FC which will be used to judge proposals. A crude but not wholly inaccurate grouping was made into major categories.

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Definitions

- Baud rate is the repetition rate on the media
- Burst transfer rate is net of any overhead introduced by encode/decode
- Performance will be based on a maximum round trip delay of 20 μsec (4 kM)
- A limit will be set for the number of passive connections allowed between two active units
- A unit is a processor, peripheral, controller, switch or any other device which complies with the physical interface and protocol
- Error rate of catastrophic events (e.g. control protocol failures) to be 10E12 on optical media, and aim for same goal with copper
- A connection is the sum of all components in the patch between the input traces on one PCB at one end to the output as PCB traces at the other

Application Area

- Known topology
- Path implicit in address
- No provision for inter-networking
- Designed for the control of peripherals
- Master-slave relationship may be dynamic
- Coordinated access to shared peripherals
- Closed system i.e. defined set of resources
- Resources dedicated to a single operating system
- Multiple processors can coexist on the same channel
- IPI and SCSI devices can coexist on the same channel
- Compatible with existing SCSI and IPI-3 command sets
- A connection will offer at least the same MTBF as a SCSI or IPI
- Connector footprint will be no larger than for an equivalent SCSI or IPI
- To fit confined spaces, cable bends no worse than equivalent SCSI or IPI

Environment

- Distance for serialized copper of 0-10M
- Office-compatible in terms of "weather," shock and vibration
- User education requirements equivalent to FDDI, and greater than for the existing copper interfaces
- Supported distance for optical transmission can vary from 5M to 2,000M (with different media and/or transmission technology); median is 500M

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Cost/Distance

- Able to use multiple technology choices
- Medium designed to accommodate future product generations
- Cost effective at 50 meters and designed to promote operation over longer distances
- Achieve approximate cost parity with a copper connection in an equivalent environment
- High performance and long distance operation possible at a cost premium that scales with the parameters

Fault Tolerance

- Able to be configured with multiple paths
- Permit online servicing of attached units
- Permit detection and isolation of failures
- Able to be configured with no single point of failure
- Allow active elements to be used if MTBF is acceptable
- Able to continue operation when units attached are powered off

Performance

- Minimize delays within attached units
- Low latency control protocol (datagram)
- Provide for control of power sequencing
- Protocol efficiency of >80% with 4K transfers
- Number of connected units supported is limited to 32
- Either sender or receiver able to throttle a transfer
- Burst transfer rate granularity of 5, 10, 50 and 200 MBs
- Flow control mechanism to reduce overhead between data bursts
- Random bit error rate of the media is a product of technology used
- Higher transfer rates may be achieved by adding parallel data paths
- Hardware error detection on boundaries (e.g. CRC on modulo 256 bytes)

Optical Extender

There is startup activity in optical interfaces going on, as evidenced by Integrated Photonics (Copper-Fiber converters), Paralan (SCSI extender) and Anil Jain (Optivision) who described his company's activities. Much of his work is Defense Agency-related, like a 4-way and 16-way crossbar switch. Del objected to the claim of a multi-way switch.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Del felt the devices were more of a multiple input/multiple output because a passive star coupler is used on each side of the switch to unite multiple fibers. The switching is point-point internally, thus preserving the full bandwidth of the fiber. The 4-way has been built and the 16-way is under development. There is a big signal loss.

Insertion Loss	(dB)	19.5	29.0
Reconfiguration	(μsec)	<20	10
Transfer Rate	(Gbs)	>1	>1
Switching Voltage	(V)	200	50

A VME bus extender is also under development. As VME is asynchronous, it is possible to have long delays due to cable length. The protocol is converted and encoded to the fiber, reducing the number of lines needed. Attachment can be transparent but there are advantages to modifying the protocol to take advantage of memory-memory transfers.

Mike Hauke (Optivision) described the implementation made of the DEC DR11-W interface, an instrumentation bus with unidirectional 16-bit data paths. It is an interesting dilemma to support this interface because only one end of the cable is defined. This has led to a wide diversity of implementations e.g. although there are Data Clocking signals, not all attached devices use the signals for clocking data.

Being an interlocked interface designed for local attachment, there is a problem supporting long distances with fiber. Mike has been successful in emulating the DR11-W but the data rate drops to 100 KBs at a kilometer, so a technique similar to SCSI Synchronous or IPI Data Streaming is used. This is known to be cause incompatibility with some DR11-W implementations.

SD3 Status

If technology does not get you, politics will. In the present environment of X3T9, politics play an important role in the progress of standards. The SD3 (permission to start a project) sailed by at the X3 (top committee) level by a vote of 36:1. The surprise was that AMP cast the negative vote.

The concerns were neither technical nor philosophical, but political. The fiasco over the FDDI connector cost AMP (and others) significant amounts of time and money. Up pops another fiber project and visions of another black hole came to mind on connectors.

The suggestion attached to the No was that "X3T9 request EIA to write the detailed specifications for the physical media components". Del questioned whether EIA would be willing to do this as the EIA work seems to be broader in scope than handling narrow application requirements.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The response was blase - it only mattered that the SD3 had been approved. There was no objection to establishing a liaison to work with EIA on the connector, as it might be useful. There is no intent to hand over a piece of the project to the whims of another committee's timetable. The FC working group will rely on EIA for technical support and assistance but EIA will not be given responsibility for the connector.

The FC group is interested in knowing more about the standards that may have been set for cabling and protection in office environments (the anticipated environment for FC installation). Wayne felt it might even be desktop now that workstations have shrunk so much in size, but the limit of 32 attached units removes the FC from this role. If there is a need to hook up large numbers of workstations it will be with a LAN, with LAN concentrators used to attach to the FC.

What About FDDI?

There is a question still hanging over the group, which is "Why does FDDI not fit the need that the Fiber Channel is intended to fill?" Paul was the one who explicitly raised the question, and it seemed like the right time to make a quick pass at listing the reasons.

- Bandwidth - at 100 Mbs, the FDDI is marginally faster than the original IPI, and half the speed of the Enhanced IPI. This is inadequate for the applications.
- Address Space - FDDI has a 48-bit address space and the Fiber Channel only needs a 5-bit space.
- Address Meaning - in FDDI the address is virtual, whereas with the Fiber Channel, the address includes the implicit connotation of the path to use
- Parallel Transmission - there is no provision in FDDI to increase bandwidth by transmitting data in parallel
- Access Delays - the defaults assumed for FDDI are for a large ring, and the smaller size of Fiber Channel requires a different philosophy
- Cost - the anticipated cost of FDDI connections is still too high for the performance achieved relative to the hoped-for Fiber Channel costs
- Station Management - a complex area the Fiber Channel will minimize

FIBER CHANNEL WORKING GROUP DECEMBER 1–2 1988

Tandem Computer hosted the meeting at its Cupertino facility.

Company	Name	Company	Name
AMD	P. Scott W. Wong	Amdahl AMP	M. Motegi C. Brill

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Company	Name	Company	Name
Ancor Comm	T. Anderson	Fujitsu America	K. Mori
AT&T	M. Kao	Gazelle μCircuits	C. Popat
	P. Puglisi	Hewlett Packard	D. Hanson
	S. Siegel	IBM	H. Brandt
Canstar	K. Malavalli		R. Soderstrom
	K. Lue Shing		H. Truestedt
CDC	F. Holland		C. Wong
	W. Sanderson	Los Alamos Labs	D. Tolmie
CDC Canada	R. Cummings	Pilkington	S. Honey
CDC Imprimis	T. Leland	Prime Computer	M. Fitzpatrick
Ciprico	B. Winterstein	Siemens	S. Van Doorn
Cray Research	E. Fromm	Tandem	A. Pauker
Data General	D. Hartig		D. Pham
ENDL	D. Allan		P. Sinykin
Fujitsu America	B. Driscal		J. Smith

Office Environment

Horst Truestedt (IBM) described the cable types that must be used for office environments.

The typical Fast SCSI, IPI or HSC cable is classified as CL2, and has to be contained inside a plenum (such as underneath a false floor). To qualify for being run without protection requires a CL2P rating. CL2P has Teflon jacketing that makes the cable quite rigid, and capable of withstanding temperatures of 2,000 degrees.

Through a fortuitous turn of fate, the building where the meeting was being held had this class of cable strung in the ceilings and there was a drop to the floor immediately behind Horst. Although CL2P cable is expensive, it is a cheaper way of wiring than having to deal with false floors or plenums.

Ron Soderstrom (IBM) covered the safety and regulatory issues associated with light emissions, be they by lasers or LEDs (Light Emitting Diodes). The FDA (Food and Drug Administration) CDRH (Center for Devices and Radiological Health) administers a relatively old law. Compliance is self-certification: a manufacturer submits documentation but never receives a notice it complies with the law. It is assumed that not getting a no means yes.

ANSI Z136-2 covers applications which would include the Fiber Channel, and the IEC (International Electrical Commission) has its own regulatory requirements in Europe.

Unfortunately, the three do not agree on what represents a safe environment. Eye damage is wavelength sensitive. The eye cannot focus well at any limit beyond 1,550 nM, there is little danger from light emissions above there.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The basic unit of measurement is the AEL (Accessible Emission Limit) which is associated with the NA (Numerical Aperture). The maximum size an eye can dilate to is 7mm, the size of a pupil. The CDRH references this to a time base of 10,000 seconds and the amount of exposure based on duty cycle of an eye. The IEC is much more restrictive as it talks of a 50mm NA. The only way to reach this size is if a magnifying glass is used to look at the signal with a naked eye.

Service and maintenance considerations are the issue, as in normal operation there would be no exposure to light. Service is fixing a problem and maintenance is likened to a user adding toner to a copier. Since maintenance is a non-issue for the Fiber Channel, service is the only consideration.

The average emission allowances are so low that some mechanical means must be designed to shut power down during service or any accidental situation when cable is exposed (such as a disconnected cable). The CDRH and IEC only cover lasers, but ANSI covers LEDs as well.

Topology

Roger Cummings (CDC Canada) described the feasible topologies for a Fiber Channel environment.

- Star
- Switching matrix
- Single rotating ring
- Dual, counter-rotating ring

A major cost item is the number of T/Rs (Transmitter/Receivers) needed for minimum operation. Don Tolmie (Los Alamos Labs) and Henry Brandt (IBM) were concerned over two T/R pairs in a unit for applications where an HSC-like simplex channel is suitable.

A single conductor ring would meet the need but every break would cause an outage. A long loop back from the last unit in a configuration would need to be reconfigured every time there was a system change. This is unrealistic, so cables have to be able to be daisy chained.

No-one could see any way to do this with only one T/R pair per unit until Schelto van Doorn (Siemens) had a brainwave.

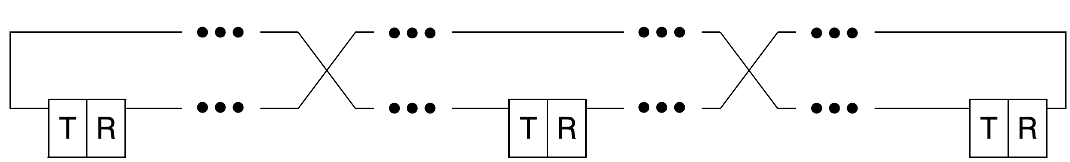

By putting a twist in the cable, so that one conductor in a dual-conductor cable is passed-through and the other is active at each connector, a daisy chain can be constructed with only one T/R pair in each unit.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

This clever technique appealed to everyone as a suitable method to use a single cable and keep a valid return path.

Principles of Operation

Data Streaming (or Synchronous SCSI) removes the speed barrier of cable turnaround delays. Control Streaming does the same thing for channel protocol but it requires a head switch (internal, not magnetic) in the way of looking at channel actions.

IPI and SCSI are connection driven protocols that do not map cleanly e.g. a typical IPI connection sequence has three turnaround delays:

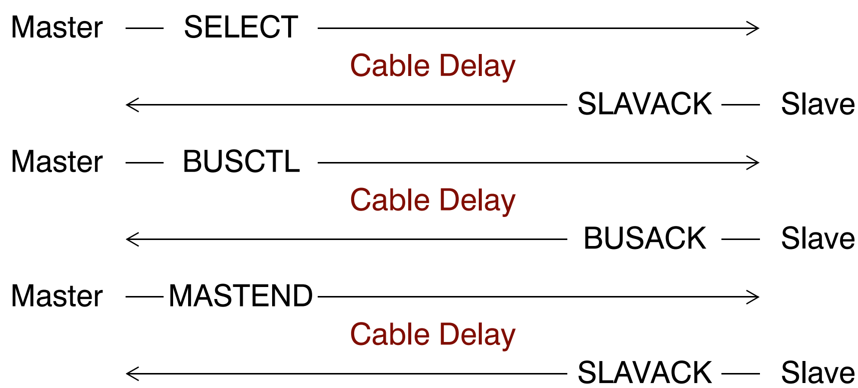

A Control Streaming protocol eliminates turnaround delays by not waiting for the acknowledgment, thereby cutting out half of the protocol steps. All the master sequences are executed at once, and some time later all the slave sequences are executed together.

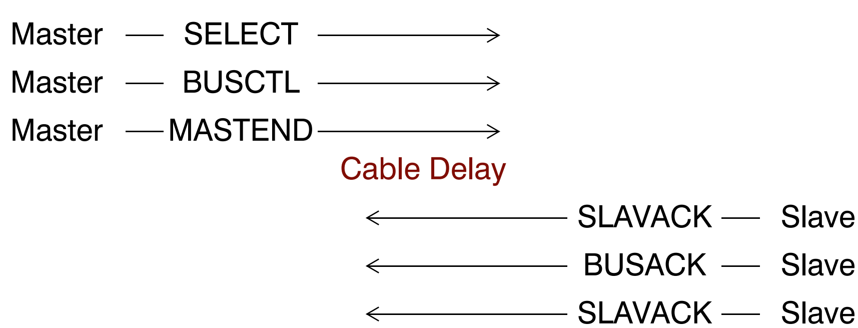

An important principle agreed upon was that only the initiator of an action on the Fiber Channel can terminate the action. This ensures that there is closure on all actions e.g. if an action by the source is not responded to by the destination then a source timeout will detect a failure. If there is a response then the source knows the action has been completed (though it need not have been successful).

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.
This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

Application

Steve Siegel (AT&T) described an implementation currently under way at Bell Labs using parallel fibers. Receiver sensitivity limits the maximum distance that can be achieved with the DC-coupled T/Rs used, and running parallel introduces skew problems, a variable of 3-10 nsec/conductor at 1 KM. There is no encoding or scrambling of data; raw data is transmitted.

Optical and electrical crosstalk from inductive pick-up is another concern in the ganged array of LEDs, and with 10-11W of power, thermal dissipation has to be properly handled to avoid a reliability problem.

The power available is approximately -18dB/M for each of the 12 LEDs in the array and receiver sensitivity is set to -30dB/M.

After accounting for optical loss considerations, approximately 7dB is left in the budget and after allowing 2dB for other losses the system margin that remains is 5 dB. Typical losses result from a wide variety of factors

Schelto emphasized the consideration that must be given to the fiber plant with regards to potential dB loss. The typically quoted loss of a mated pair connector is very optimistic at 0.5-0.7dB. This value is after cutting the fiber, connectorizing, and then measuring; and produces significantly better results than in the real world where different fibers are intermated using connections made at different times.

Unless care is taken there can be serious loss problems. In situations where Siecor has installed fiber plant using mixed connectors there has been such severe degradation they have had to be recabled or reconnectorized. The only way to ensure the least amount of loss in a fiber plant is to use a single vendor throughout for both fiber and connectors.

Steve feels that the maximum that can be expected from LEDs with DC-coupled T/Rs is 250 Mbs over up to 1 KM. This would appear to meet most application needs, but not for Henry Brandt. Henry wants to run at 200 MBs for up to 10 KM in order to execute remote backups. This could either be done using lasers and single mode fiber or via repeater stations.

Dynamic skew management by an elastic store or tapped delay lines must be used to reduce skew between the individual conductors in a parallel cable. AT&T is investigating chemical composition control as a way to achieve the same refraction in each conductor of a parallel cable.

Steve's work is still in the development phase. In response to a question on how much information can be provided, Phil Puglisi (AT&T) indicated he is more than willing to report on test results for environments of interest to the working group.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The different performance characteristics of parallel cable are priced on a simple ratio of twice as much for each improvement.

- $ 3.25/M – unsheathed with 7 nsec skew
- $ 7.50/M – sheathed with 7 nsec skew
- $ 15.00/M – sheathed with 3 nsec skew

Wayne Sanderson (Control Data) did some quick calculations on the cost of a fiber plant strung with minimum skew cable and asked about the tradeoffs in a comparison of different technologies, such as running bit serial at 1 GBs where the cost is in T/Rs and the signal processing.

Running parallel with only one clock requires expensive fiber and when long distances are involved it gets expensive. The $15,000 for 1 KM of 3 nsec skew cable will "*buy a lot of silicon*." It might be significantly cheaper to transmit data with embedded clocks over unmatched parallel fibers and use LSI to compose and decompose the data as well as handle the FIFO buffers.

HSC

Don Tolmie described the impetus behind the campaign for the HSC (High Speed Channel). Karl Heinz Winkler is a German scientist who was convinced that moving pictures would help identify and solve problems in his research.

A configuration was "hobbled together" (at a cost of about $1M) to verify this conviction. The objective was to demonstrate and measure the airflow disturbances which are the result of physical motion by an object.

Winkler discovered several analytical flaws from the animation of 5-year old data. Watching the screen he was able to identify both a physics problem and a numerical instability in the simulation. These results led to a wave of enthusiasm for pictorial representation of formula and data.

Computers are good at crunching numbers but not at displaying images.

People are image oriented, and are poor at interpreting pages of numbers. A new level of technology has to be developed to create movement without any glitches. The eyeball and mind naturally focused on glitches and can miss the real problem.

Don had a video tape with him to illustrate the results being achieved with existing hardware and software. As it takes about 10 hours of Cray time to develop sufficient data for 30 seconds of motion, it may have cost even more than a Michael Jackson MTV spot.

Another problem raised by motion display is how and where to store the data generated so that scientists can replay the screen over and over again.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

Don's perception of HSC is as a big circuit switch whereby several channels of data can be routed over a matrix switching net to simultaneously service a myriad of users. As this is extremely expensive, he is looking for another and more cost effective technique to achieve distance.

Principles

The Fiber Channel will replace the first two layers of the SCSI, IPI and HSC physical interfaces.

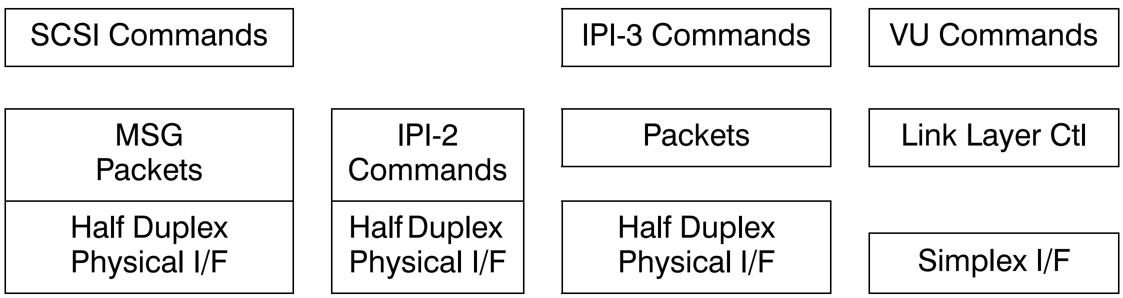

Level definitions similar to those of IPI are used, with FC-0 defining the cabling medium, connectors and transmission criteria and FC-1 defining the signaling protocol. The technology used in FC-0 will dictate the performance achievable per conductor, and more than one conductor can be ganged together to build a transfer rate many times above the speed of a minimum attachment.

There are doubts over whether IPI-2 really belongs in the list of intended support, but all options are being left open at this time. At present there is no reason to eliminate IPI-2 as a command set that will be supported. HSC users have Vendor Unique commands. This correlates with the objective of ensuring that it will be easy for proprietary channels to be mapped across to the Fiber Channel.

FC-1 will be consistent across a range of FC-0 alternatives. Transmission criteria such as framing and the data reduction scheme are not considered to be part of the FC-1 signaling protocol but belong to FC-0.

This allows the best technique, be it encoding, scrambling or naked transmission with in-line EDAC (Error Detection And Correction), to be transparent to supporting logic (which can then be common across all implementations).

Transmission is serialized but that does not mean only one conductor. The Primary Bus Group has the least number of conductors needed to carry the signaling protocol plus one Data Bus Group (the minimum is one conductor). Additional Data Bus Groups may be added.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.
This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

The purpose of additional Data Bus Groups is to increase transfer rate, and it need not be symmetric i.e. if an attached peripheral is a high resolution graphics terminal then there might be two Primary Groups for Input/Output but seven output Data Bus Groups to multiply the output bandwidth to 8X.

If an application has non-coincident periods of heavy traffic, then all the Data Bus Groups might be input at one time and then be output at another. The control for the dynamic allocation of Data Bus Groups is not going to be specified in the standard.

In rings with limited reliability, single attachment star configurations, or point-point (HSC-like) configurations, there need only be one Primary Bus as the loop around satisfies the criterion that for every In there be an Out.

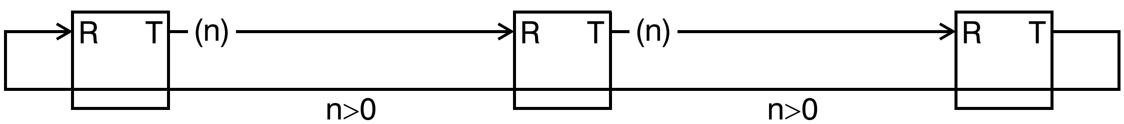

In other configurations, a Fiber Channel would have two Primary Buses.

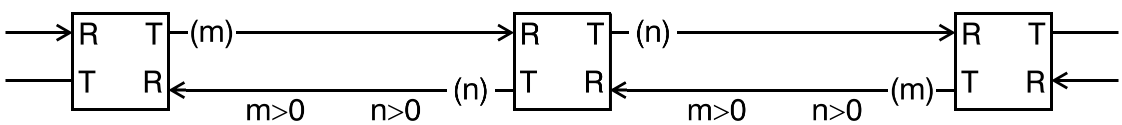

Conductors may be either fiber or copper, because copper is much easier to connectorize and use intra-cabinet. A parallel copper channel could be laid down as a backplane to plug peripherals into, and there is a hope that it will be an attractive alternative physical interface for SCSI-3.

The estimated reduction in cabling count for a parallel implementation vs a serialized is expected to be on the order of 25:1 i.e. a 2-conductor coaxial or fiber implementation would be equivalent to a 50-pin parallel version.

All data transmissions will either be short control sequences (which may be called tokens) or bursts. Every burst will have a prefix that includes the command interface (IPI/SCSI/other), size, source and destination addresses, and its own unique identification number. A packet consists of one or more bursts, and bursts from several packets may be interleaved, kept separate by the Burst Prefix identification.

Frank Holland (CDC) pleaded for the Burst Prefix to have at least a 32-bit multiple, preferably 64-bit, to provide for future system memory addressing growth. A Link bit opens up the possibility of being able to route data from one source to more than one destination.

The Burst Suffix will contain integrity check data for a polynomial that is known to both the source and destination. If zero, then there is no check and this capability will be in addition to anything which may already be in the FC-0, or it may be an encryption check for security.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.

This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!

Working Documents

The documentation on the Fiber Channel is expected to become quite bulky and the working group is about to enter the development phase as this meeting marked the end of the learning and awareness period.

Although proposals had been solicited at the last working group, none were forthcoming, despite several known to be under implementation. All of them are being kept proprietary because of management concerns over the release of information prior to product announcement. However, in the event that any one company breaks this barrier, it is more than possible several others will be brought forward.

In the interim, it was agreed that a significant amount of information is already in the public domain, such as the description of MacroLAN (provided by ICL) plus the Enhanced IPI, SCSI-2 and HSC standards. Even in the absence of proposals there is sufficient material to advance the Fiber Channel.

ACRONYM GLOSSARY

ADPCM Adaptive Differential Pulse Code Modulation

AEL Accessible Emission Limit

ANSI American National Standards Institute

APD P-N junction with Avalanche Multiplication

ASC Accredited Standards Committee

BER Bit Error Rate

CD Compact Disk

CDC Control Data Corporation

CDRH Center for Devices and Radiological Health

CE Customer Engineering

CPU Central Processing Unit

CRC Cyclic Redundancy Check

DAT Digital Audio Tape

DC Direct Current

DQDB Distributed Queueing Dual Bus

ECC Error Correction Code

ECL Emitter Coupled Logic

ECMA European Computer Manufacturers Association

EDAC Error Detection And Correction

EIA Electronic Industries Alliance

ELED Edge Light Emitting Diode

EMI Electro Magnetic Interference

FC Fibre Channel

FC Ferrule Connector

FDA Food and Drug Administration

FDDI Fibre Distributed Data Interface

FIFO First In First Out

FIPI Fiber IPI

FL Fiber Optic Link

Special Edition

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

FO	Fiber Optic
HP-FL	Hewlett Packard Fiber Optic Link
HSC	High Speed Channel
I/F	InterFace
ID	Inner Diameter
IEC	International Electrotechnical Commission
IEEE	Institute for Electrical and Electronics Engineers
IPI	Intelligent Peripheral Interface
IPI-2	Intelligent Peripheral Interface
IPI-3	Intelligent Peripheral Interface
ISDN	Integrated Services Digital Network
ISO	Organization for International Standards
LAN	Local Area Network
LED	Light Emitting Diode
LLC	Link Layer Control
MAC	Media Access Control
MAN	Metropolitan Area Network
MIPS	Millions of Instructions Per Second
MSG	Message
MTBF	Mean Time Between Failure
MTV	Music Television
NA	Numerical Aperture
NASA	National Aeronautics and Space Administration
NIU	Network Interface Unit
OD	Outer Diameter
OSI	Open Systems Interconnect
PCB	Printed Circuit Board
PIN	P-N junction with Intrinsic region
PMD	Physical Media Dependent
PMMA	PolyMethyl MethAcrylate
POTS	Plain Old Telephone Service
QPSX	Queued Packet and Synchronous Switch
RFI	Radio Frequency Interference
RISC	Reduced Instruction Set Computer
ROM	Read Only Memory
SA	Shipboard Applications
SAFENET	Survivable, Adaptable Fiberoptic NETwork
SC	Subscriber Connector
SCSI-2	Small Computer System Interface
SCSI-3	Small Computer System Interface
SD3	Project Request
SLED	Surface Light Emitting Diode
SONET	Synchronous Optical NETwork
SSD	Solid State Disk
TAB	Technical Advisory Board
TV	Television
VME	Virtual Machine Environment
VU	Vendor Unique
WAN	Wide Area Network
XTP	eXpress Transfer Protocol

A license is granted to FCIA member companies; this ENDL Letter Special Edition may be copied for wide distribution.
This ENDL Letter Special Happenings is to be shared with all FCIA members and T11 friends. Congratulations T11!