FC signal specification architecture – application to higher speeds

Bill Ham, HP

(Not an official T11 or HP position statement)

FC signal specification architecture history

- FC developed a basic link performance requirements specification methodology several years ago that was documented in a technical report called MJS (Methodologies for Jitter Specification – now replaced by MJSQ), and in three standards: FC-PH3, FC-PI, and FC-PI-2
- The same general methodology is used by parallel SCSI standards, ATA, SATA, and SAS: the link performance requirements apply at the device connectors (not at the chips) – This is called the connector based method in this presentation
- Other technologies (e.g., Ethernet XAUI, OIF backplanes) use a different methodology where the chip to chip ‘channel’ (that may or may not contain connectors) is defined and the link requirements are specified basically at the chip pins – This is called the channel based method in this presentation
- These two methodologies are similar in some ways and different in others

Connector based /channel based

- The connector based method is best suited for applications where the important standard interoperable interface is at a device, the channel based method is best suited for applications where the channel is known and controlled by the product. For example, a channel optimized application could be in a blade based architecture where the backplane and blades are all part of the same system design
- While SAS has desirable implementations where the channel based method could be applied (such as in a blade server environment for example) most applications involve connecting storage devices such as HDD’s to backplanes or cables where the connector based method is significantly preferred if only one method is used
- The SAS architecture assumes independent link components that separate at the connectors and this needs to be preserved in future variants
- The channel based method assumes that a common chip interface is the driving force for the technology.
- While the two methods are not mutually exclusive, specifying both chip and connector requirements is viewed by many as excessively restrictive – chip specifications are not contained in FC, SAS, or parallel SCSI
- Channel based methodology that uses a ‘golden’ channel is useful for emulating the frequency dependent loss properties of the interconnect in transmitter device compliance testing

Link design vs Link component performance specifications at connectors

- When new technologies are being developed it is common for the developers to concentrate on the conditions that allow a complete link to operate – this is intrinsically an overall link design activity (and is also one of the necessary conditions for the link and is hugely useful in determining the basic properties needed)
- The requirement of having an open specification where the components each are independently specified intrinsically means dealing with the performance at the separable connectors and these requirements do not come directly from the overall link design – it is a separate piece of work to divide the budgets between the link components (and to consider the interactions between the components when connected together)
- Ultimately the requirement is for the link to operate with a transmitter device, an interconnect, and a receiver device that each have performance requirements that are independent of each other
- So, unless there are no connectors in the link where interoperability is desired, an interactive process between both (1) the total link functioning and (2) the performance requirements at the connectors is needed

Some properties of the connector based methodologies

- All signal quality specifications apply at defined points around separable connectors where the system comes apart
  - This is good for folks who make higher level components like HDD’s, HBA’s, switches, raid controllers, JBOD’s, cable assemblies, backplanes with connectors, etc and creates interoperability at the system component level
  - This is not as good for folks who make chips, connectors, bulk cable and other lower level components since there are no specifications that directly apply to these components – specifications for these components are part of the design for the higher level components
  - The connector based methodology avoids the trap of overspecification that leads to real links operating much too conservatively

Some properties of the connector based methodologies

- Connector based methodology does not attempt to dictate how components are designed (except for the mating interfaces of the connectors) by forcing specific requirements like wire gauge or CMOS technology – rather:
  - Two classes of signal performance are specified: signal output and signal tolerance
  - Signal output specifications apply to signals coming out of an interoperability point into a standard load
  - Signal tolerance specifications apply to the ability of the downstream portion of the link to deliver adequate BER with a specified worst case signal launched into the interoperability point from an ideal source
- Components that meet both the signal output requirements and the signal tolerance requirements may be designed any way that accomplishes compliance with these requirements

Other properties of the basic FC/SAS methodologies

- FC and SAS require signal performance measurements to be done in a way that relates the signal specifications to the link BER performance
  - Link BER is only visible after the link receiver device has detected all the bits from the signal – the signal itself does not have a BER property
  - Signal measurement methodologies emulate certain assumed properties of the link receiver in terms of frequency tracking dynamics and response to data pattern changes
  - One should still expect there to be a gap between signal measurements and observed link BER performance because the signal cannot be measured in a way that closely follows the properties of the specific receiver being used in the link <u>unless</u> the properties of the specific receiver in the link are known in detail
  - Unless the link receiver device is very weak (i.e., barely compliant) it is to be expected that the observed link BER performance will be better than suggested by a signal measurement alone

Signal performance requirements at separable connectors

Name	Symbol	Description
Differential signal output	DSO	Jitter eye measurement (Note 1) with defined data pattern out of a laboratory grade electrical load - signal measured through the mated connector used in service Measures the performance of all upstream portions of the link
Differential signal tolerance	DST (Note 2) (Note 3)	BER measurement with a defined data pattern using a signal generated from laboratory grade calibration and launch conditions - signal calibrated through a laboratory grade mated connector Measures the performance of all downstream portions of the link
Differential upstream return loss (Note 4)	SDD22	Return loss measurement looking upstream using differential S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all upstream portions of the link
Differential downstream return loss (Note 4)	SDD11	Return loss measurement looking downstream using differential S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all downstream portions of the link
Common mode signal output	CSO	Jitter eye measurement (Note 1) with defined data pattern out of a laboratory grade electrical load - signal measured through the mated connector used in service Measures the performance of all upstream portions of the link
Common mode signal tolerance	CST (Note 2 (Note 3))	BER measurement with a defined data pattern using a signal generated from laboratory grade calibration and launch conditions - signal calibrated through a laboratory grade mated connector Measures the performance of all downstream portions of the link
Common mode upstream return loss (Note 4)	SCC22	Return loss measurement looking upstream using common mode S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all upstream portions of the link
Common mode downstream return loss (Note 4)	SCC11	Return loss measurement looking downstream using common mode S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all downstream portions of the link

Note 1 - see MJSQ for details on jitter eye measurements
Note 2 - Signal tolerance measurements for differential and common mode properties cannot be separated because the result of a tolerance measurement is the BER for the link. The signal used for signal tolerance measurements contains the worst case combination of differential, common mode, and data pattern properties.
The differential and common mode content in the signal used for signal tolerance is the specified quantity in this table.
Note 3 - Signal tolerance methods are described in more detail in MJSQ where four kinds of jitter content are described in these signals.
Note 4 - Use of S parameters assumes that the relevant portions of the links behave linearly - such behavior may not always exist in active devices or where magnetic coupling elements are used - the active devices are set to their nominal bias conditions during this measurement to minimize the impact of non-linear properties.

Inter-enclosure Environment (Gamma Points)

Possible Configurations (Notes 1, 2, 3, 4)

Enclosure Boundary

Alpha	-------	-------	Gamma
Alpha	-------	Delta	Gamma
Alpha	Beta	-------	Gamma
Alpha	Beta	Delta	Gamma

Connection of any Gamma point to any other Gamma point Is allowed

Enclosure Boundary

Gamma	-------	-------	Alpha
Gamma	Delta	-------	Alpha
Gamma	-------	Beta	Alpha
Gamma	Delta	Beta	Alpha

Extended Intra-enclosure Environment (2 Enclosures - No Gamma Points)

(If no Gamma point exists the environment remains intra enclosure even if the Beta or Delta points are in different enclosures – shielded interconnect assumed)

Possible Configurations (Notes 2, 3, 4)

Alpha	-------	Delta	-------
Alpha	Beta	-------	-------
Alpha	Beta	Delta	-------
Alpha	-------	-------	-------

Connection between enclosures is allowed (repeaters required for some connections)

-------	Delta	-------	Alpha
-------	-------	Beta	Alpha
-------	Delta	Beta	Alpha
-------	-------	-------	Alpha

Intra-enclosure Environment

Possible Configurations:
(Notes 2, 3, 4)

Alpha - Alpha
Alpha - Beta - Alpha
Alpha - Beta - Beta - Alpha

No connection between enclosures

Intra-enclosure Environment

Possible Configurations:
(Notes 2, 3, 4)

Alpha - Alpha
Alpha - Beta - Alpha
Alpha - Beta - Beta - Alpha

Note 1: Repeaters are required in the enclosure when the enclosure includes both Beta and Gamma points in the same link -- Repeaters preserve independent amplitude budgets for both intra and inter environments. If Retimers are used to provide this function, independent jitter budgets are also preserved.

Note 2: Signal requirements for Alpha points associated with Beta points or intra-enclosure Alpha to Alpha configurations may be different from the signal requirements for Alpha points associated with Delta or Gamma points. No specifications are given for Alpha points in FC-PI-2. Alpha points only exist with enclosures

Note 3: As required by the application, a Retimer may be inserted at any interoperability point in a configuration for purposes of compliance conversion to any other interoperability point.

Note 4: The configuration on the left is independent of that on the right and vice versa

FCIA slide (www.fibrechannel.org)

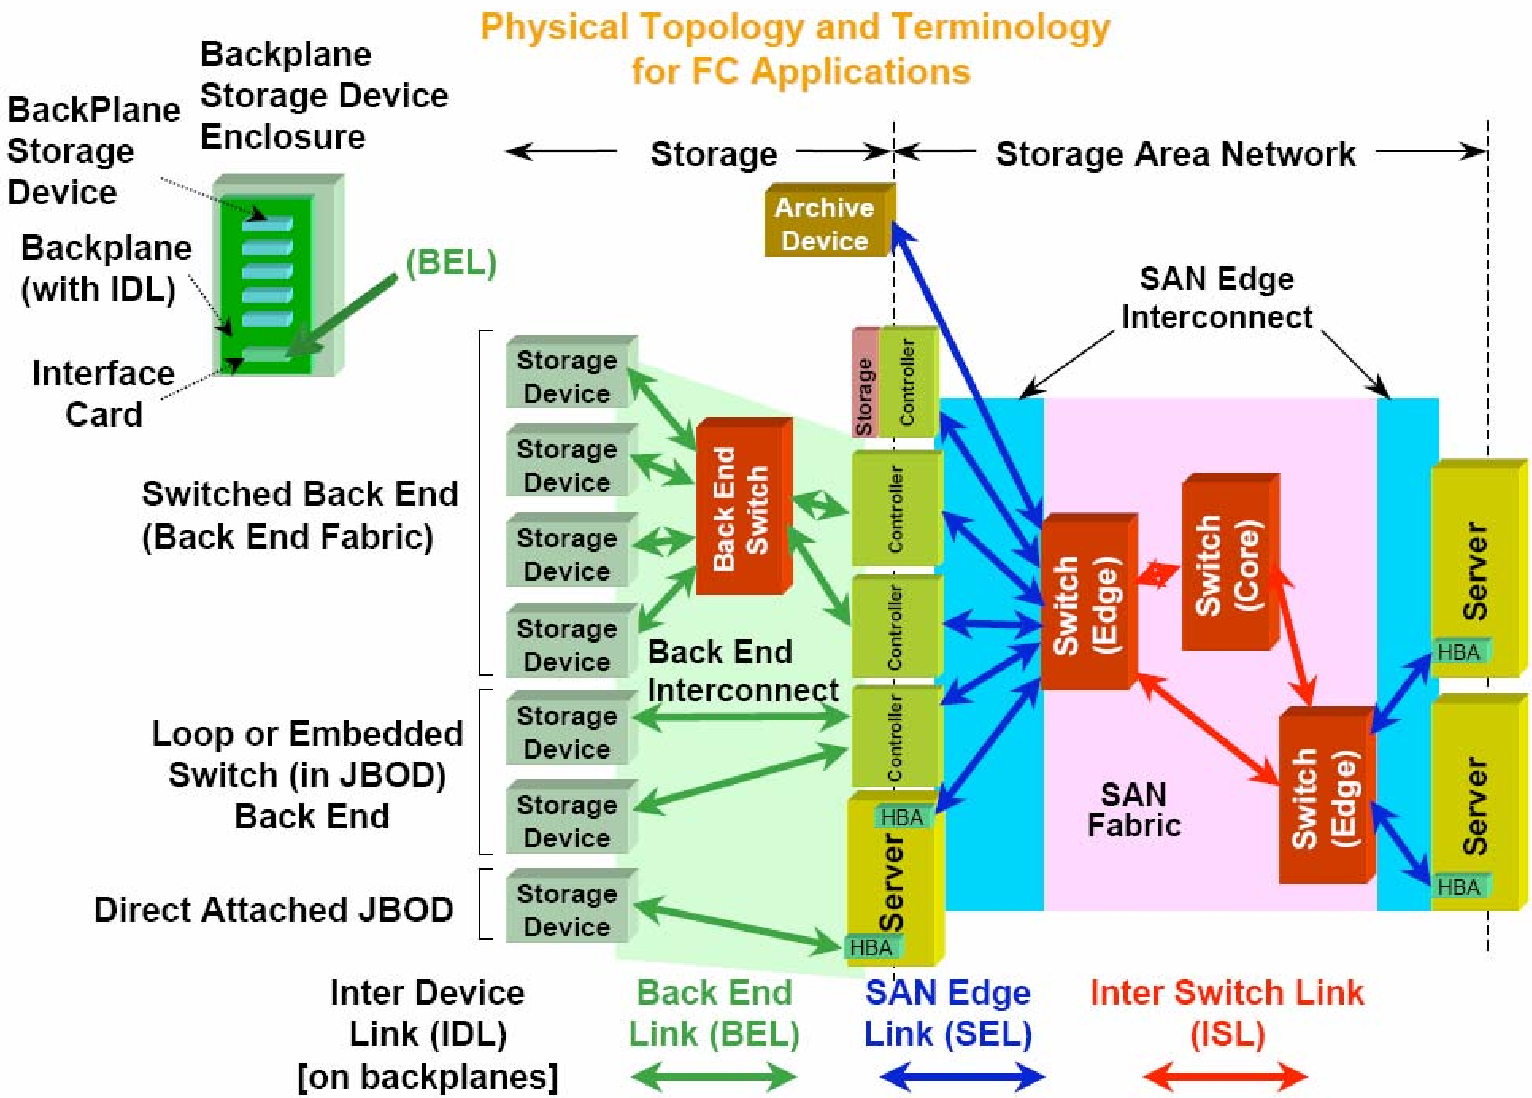

FCIA slide (www.fibrechannel.org)

Fibre Channel Infrastructure Application Matrix

Market	Connection	Length
Metro [Optical]	ISL	>= 5km
Multi Building (campus)[Optical]	ISL	300m - 5km
Single Building (Local) [Optical]	ISL	30m - 5km
Datacenter or Rack[Optical]	ISL/SEL/BEL	0m - 100m
Datacenter or Rack [Copper]	ISL/SEL/BEL	0m - 15m
Backplane [Copper]	IDL	0.6m

FCIA slide (www.fibrechannel.org)

Fibre Channel Speed Chart

Base2*

Product Naming	Throughput (MBps)	Line Rate (Gbaud)	T11 Spec Completed (Year)	Market Avaiability (Year)
1GFC	200	1.065	1996	1997
2GFC	400	2.125	2000	2001
4GFC	800	4.25	2003	2005
8GFC	1,600	8.5	2006	2008
16GFC	3200	17	2009	2011
32GFC	6400	34	2012	Market Demand
64GFC	12800	68	2016	Market Demand
128GFC	25600	136	2020	Market Demand

Base10**

10GFC	2400	10.52	2003	2004

*Base2 used throughout all applications for Fibre Channel Infrastructure and devices. Each speed maintains backward compatibility at least two previous generations (I.e., 4GFC backward compatible to 2GFC and 1GFC)

**Base10 commonly used for ISL, core connections, and other high speed applications demanding maximum bandwidth

Lines Rate: All speeds are single-lane serial stream
Dates: Future dates estimated

General view of 8GFC

- 8GFC should be the same as slower speeds wherever possible: backwards compability is assumed and required for at least two slower speeds
- This means that the same port may be capable of operating at different speeds by only changing the properties of the electronics in the port: the cable plant does not change
- This precludes the use of multilane variants to go faster and offers the possibility of significant cost and complexity savings
- Notice that the application matrix does not say that the lengths are different just because the speed is different
- 8 GFC will be specified in FC-PI-4, Link Equalization Enhanced Variants (LEEV): it is assumed that equalization may be required to achieve the required distances
- Only the minimum required equalization will be in the standard
- Equalization may be in the transmitter device, in the interconnect, or in the receiver device
- While receiver device equalization may be used it is not desirable because it requires complicating the signal measurements with standard reference receivers
- Adaptive equalization is also possible but not desirable because it complicates the signal measurements

And about connectors…

- It is periodically fashionable to blame connectors for limiting the speed attainable
- As far as the standards are concerned it is only the mating interface of the connectors that needs to remain backward compatible: the rest of the connector and its mounting interface may need to change, however – the SCA-2 is good at least to 8.5G (SCA-2’s are used mostly for Beta points)
- In most cases it is not the mating interface that causes the serious issues: try the mounting on the board

FC is also considering new variants that can use Category 5, 6, 7 cable plants

- Category cable plants (both shielded and unshielded) are installed in many places for Ethernet LAN applications
- FC is presently investigating the requirements and desirability of introducing a new (but not compatible with the present FC installed base) set of variants for 1G, 2G, 4G, 8G and 10G
- If it proves desirable there could be new lower costs associated with FC technology

And speaking of cost….

- Cost shall not increase when the speed increases! (although there may be a small bump in the early introduction stages)
- Maintaining backward compatibility goes a long way towards managing costs because one can introduce a product that is capable of higher speed without waiting for the rest of the world to catch up – just use it at the lower speeds until the rest of the system is ready for the higher speed

Summary

- FC methodologies are a nearly perfect match to the requirements of SAS
- They also set SAS up for optical if the need should arise
- The biggest differences are the FC use of single serial stream (no multilane), not using OOB, and a somewhat higher data rate
- Changes like different encoding, adaptive equalization, exotic materials, and new connector interfaces are being resisted until there is no other choice