Serial Attached SCSI
Physical layer

by Rob Elliott
HP Industry Standard Servers
Server Storage Advanced Technology
elliott@hp.com http://www.hp.com
30 September 2003

Notice

- These slides are freely distributed by HP through the SCSI Trade Association (http://www.scsita.org)
- STA members are welcome to borrow any number of the slides (in whole or in part) for other presentations, provided credit is given to the SCSI Trade Association and HP
- This compilation is © 2003 Hewlett-Packard Corporation

SAS standard layering

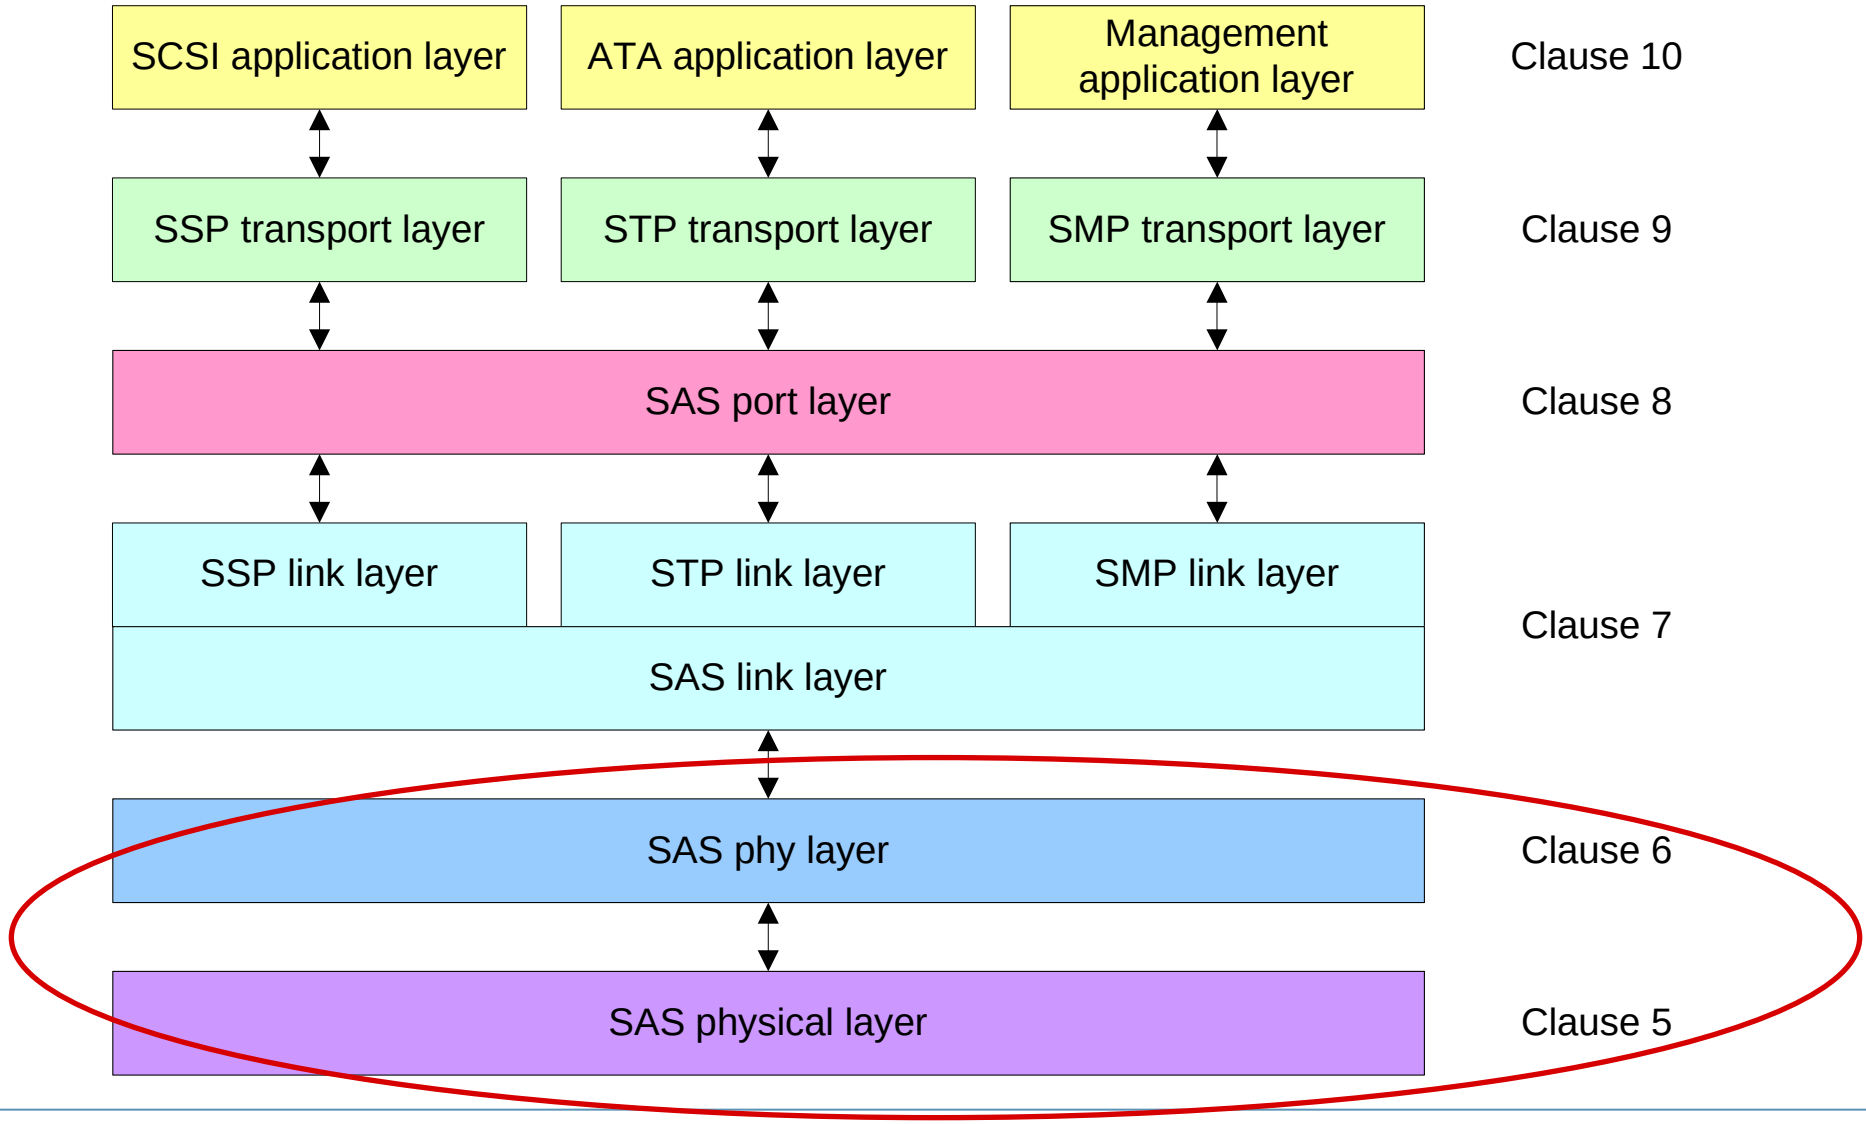

SATA 1.0a standard layering

- For SATA 1.0a from the private Serial ATA working group

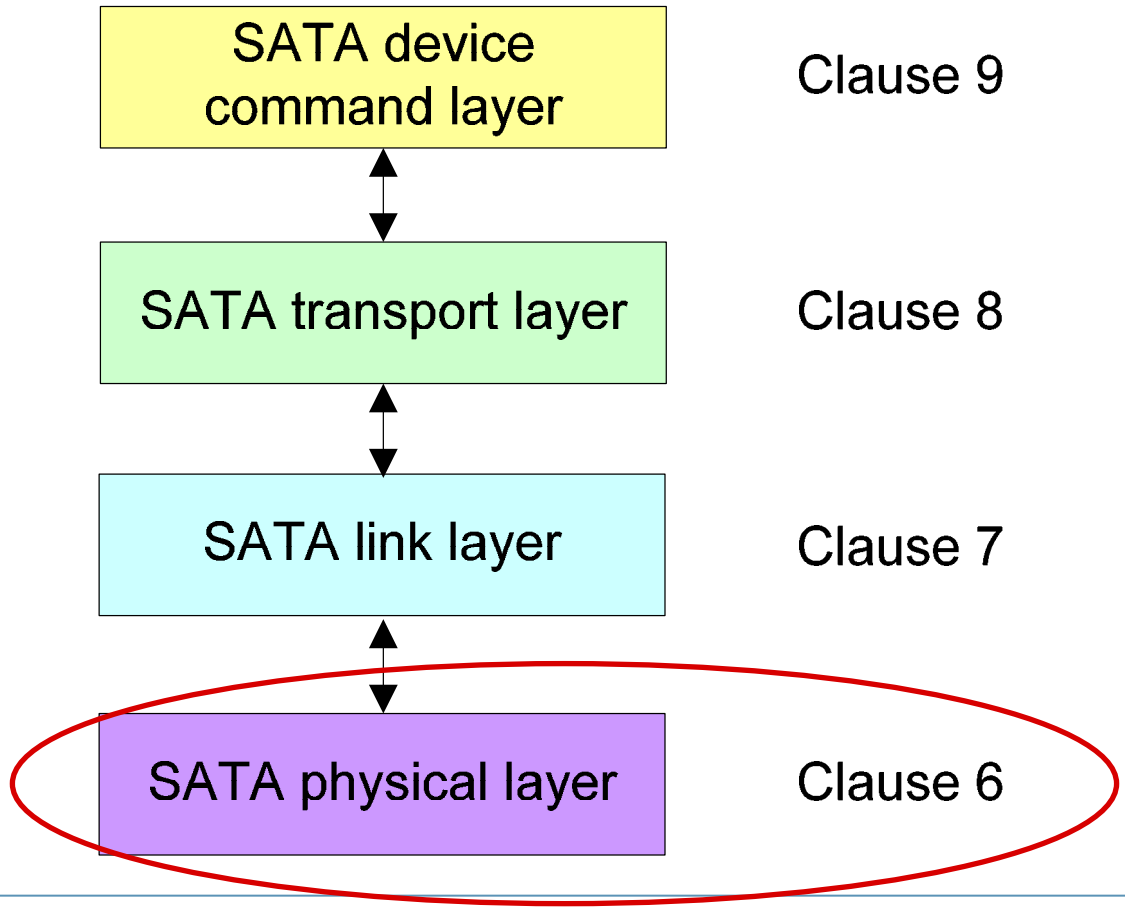

ATA / ATAPI-7 standard layering

- For the public standard ATA / ATAPI-7
- Subject to change by T13 standards committee

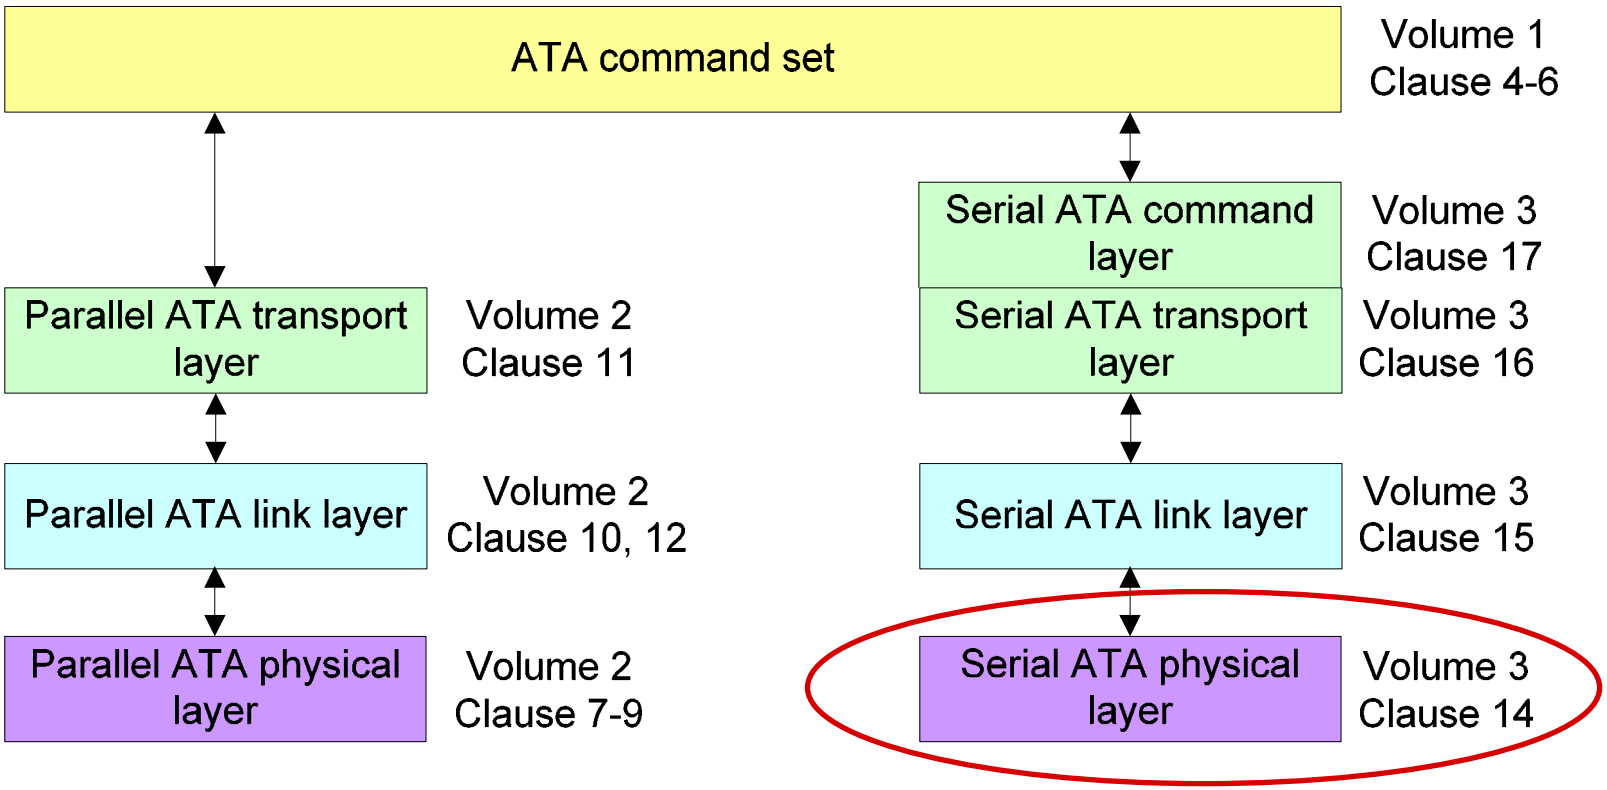

SAS clause 5 – Physical layer

- Differential signaling
- Passive interconnect
  - Internal cables / connectors
  - External cables / connectors
- Electrical characteristics
  - SATA
  - SAS
  - READY LED signal
  - Preemphasis and equalization

Differential signaling

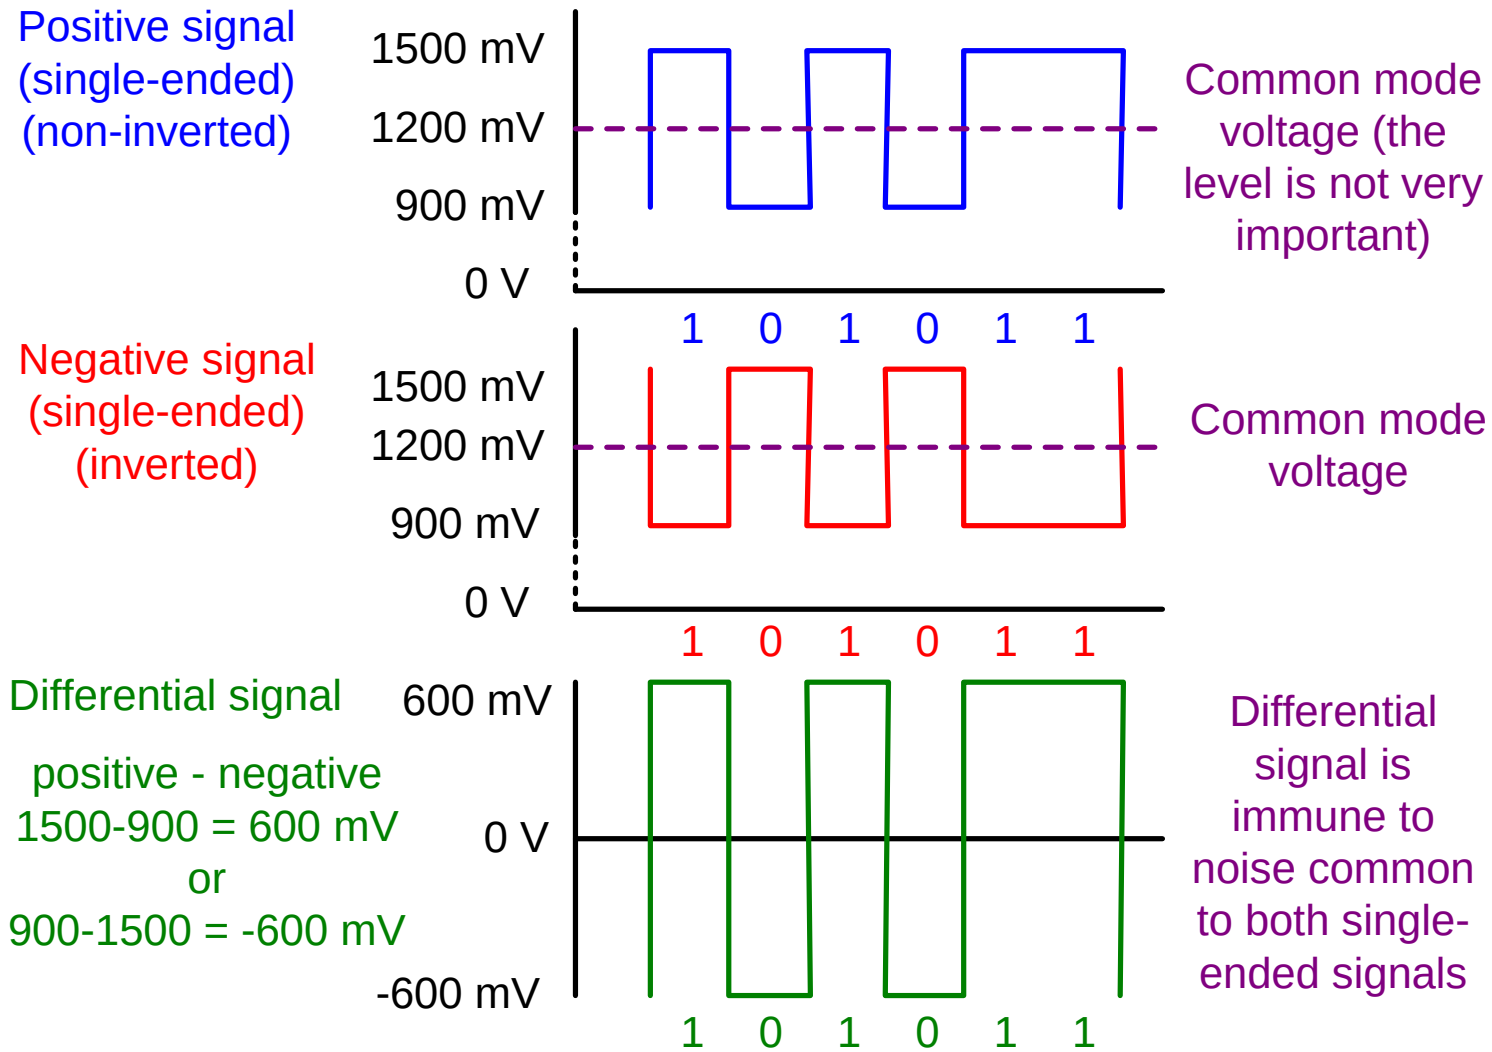

Physical layer - Passive interconnect (Cables and connectors)

SATA cables and connectors

- SATA 1.0 was defined for internal use only (e.g. inside a PC)
- 1 m internal cable
- No external connectors / cables

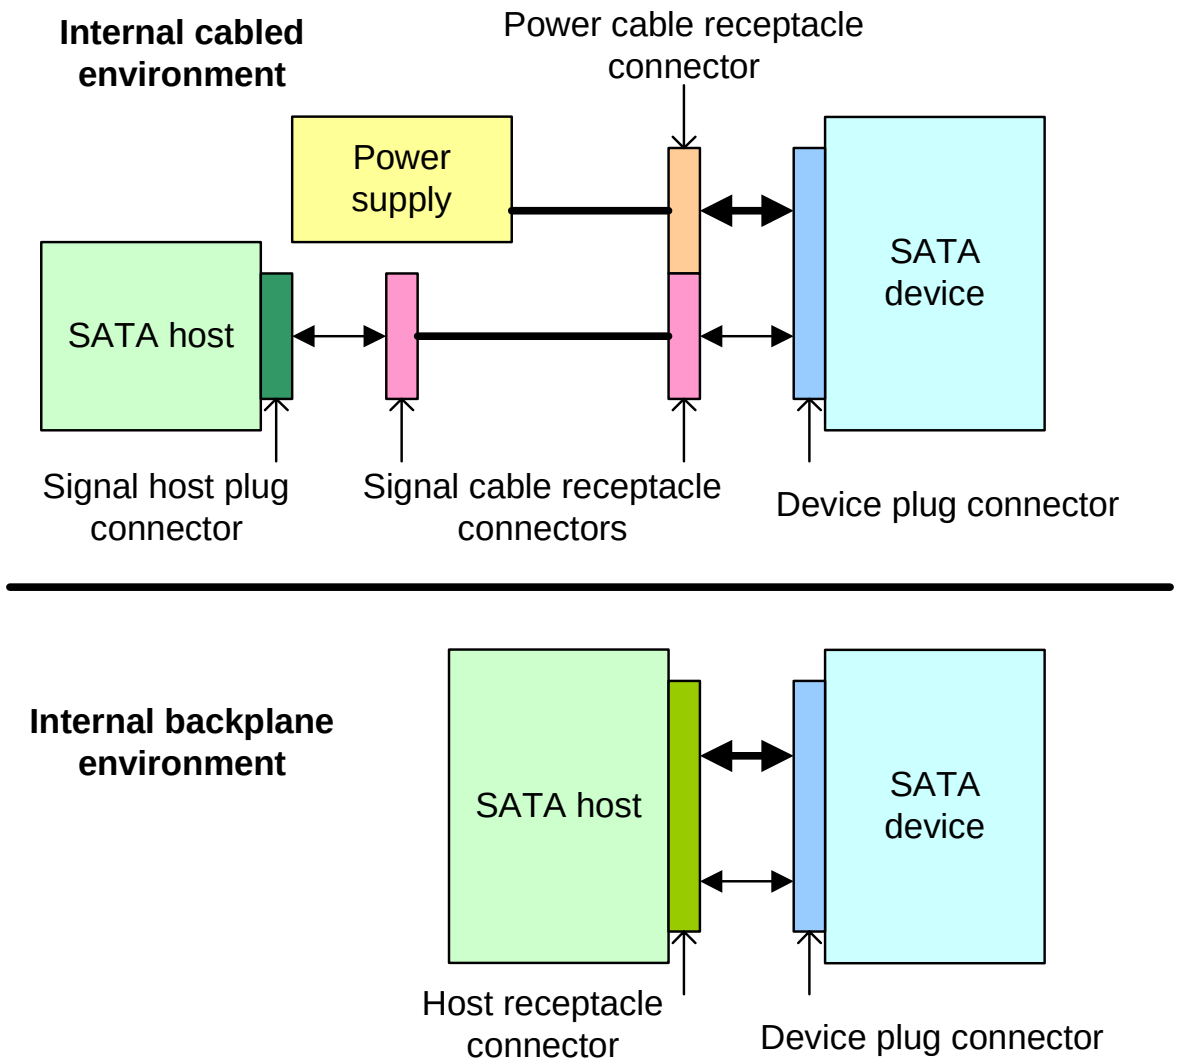

SATA device connector

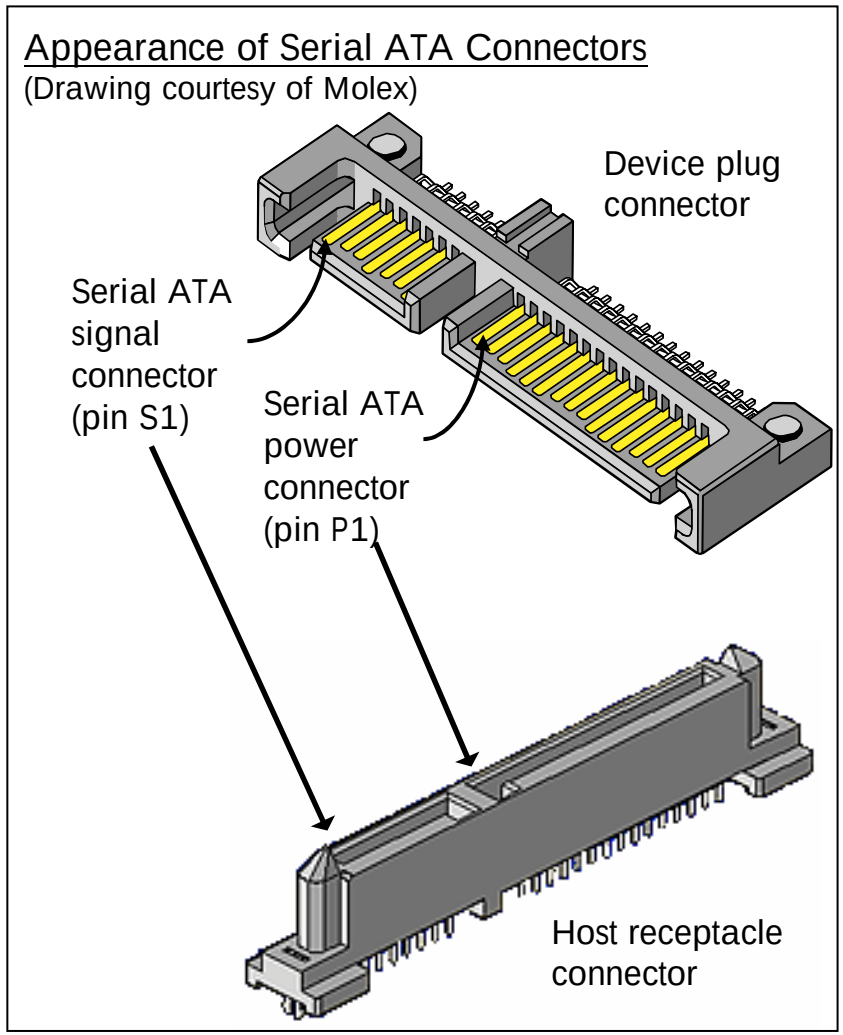

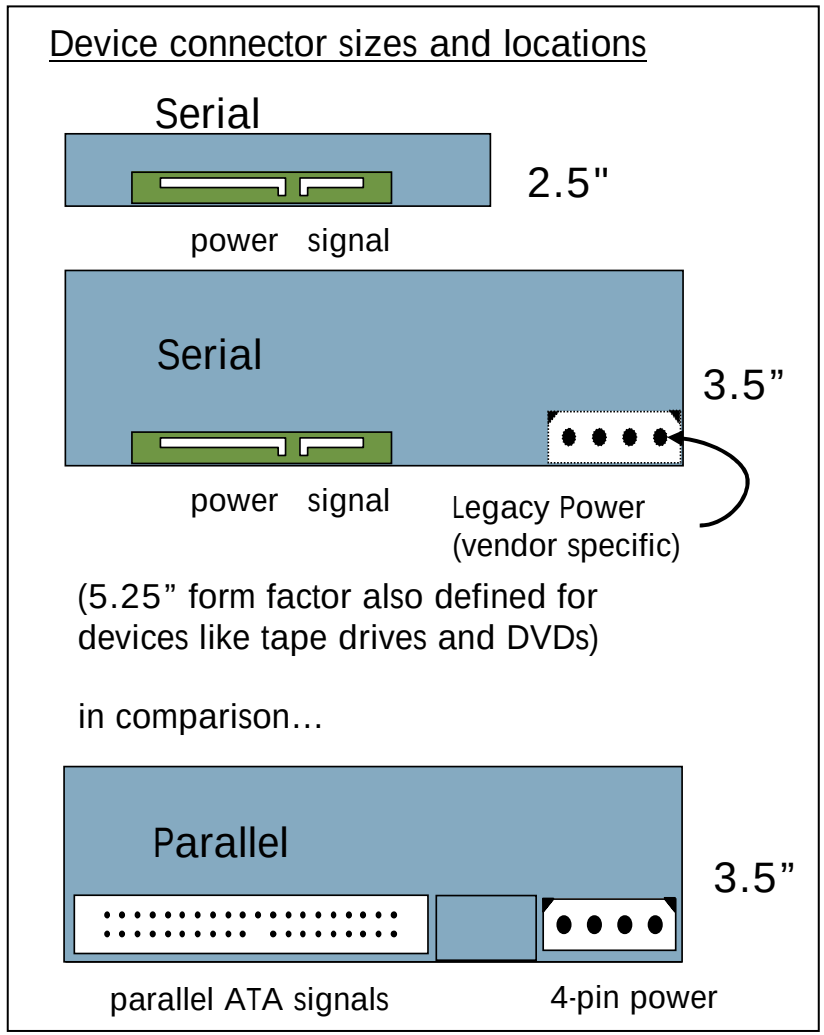

SAS cables and connectors - internal

- Similarly, SAS defines internal environments
- Backplanes support two physical links
- Cables mainly support one phy (for debug, two physical link cables might be useful)
- HP proposing 4x internal connector for SAS-1.1

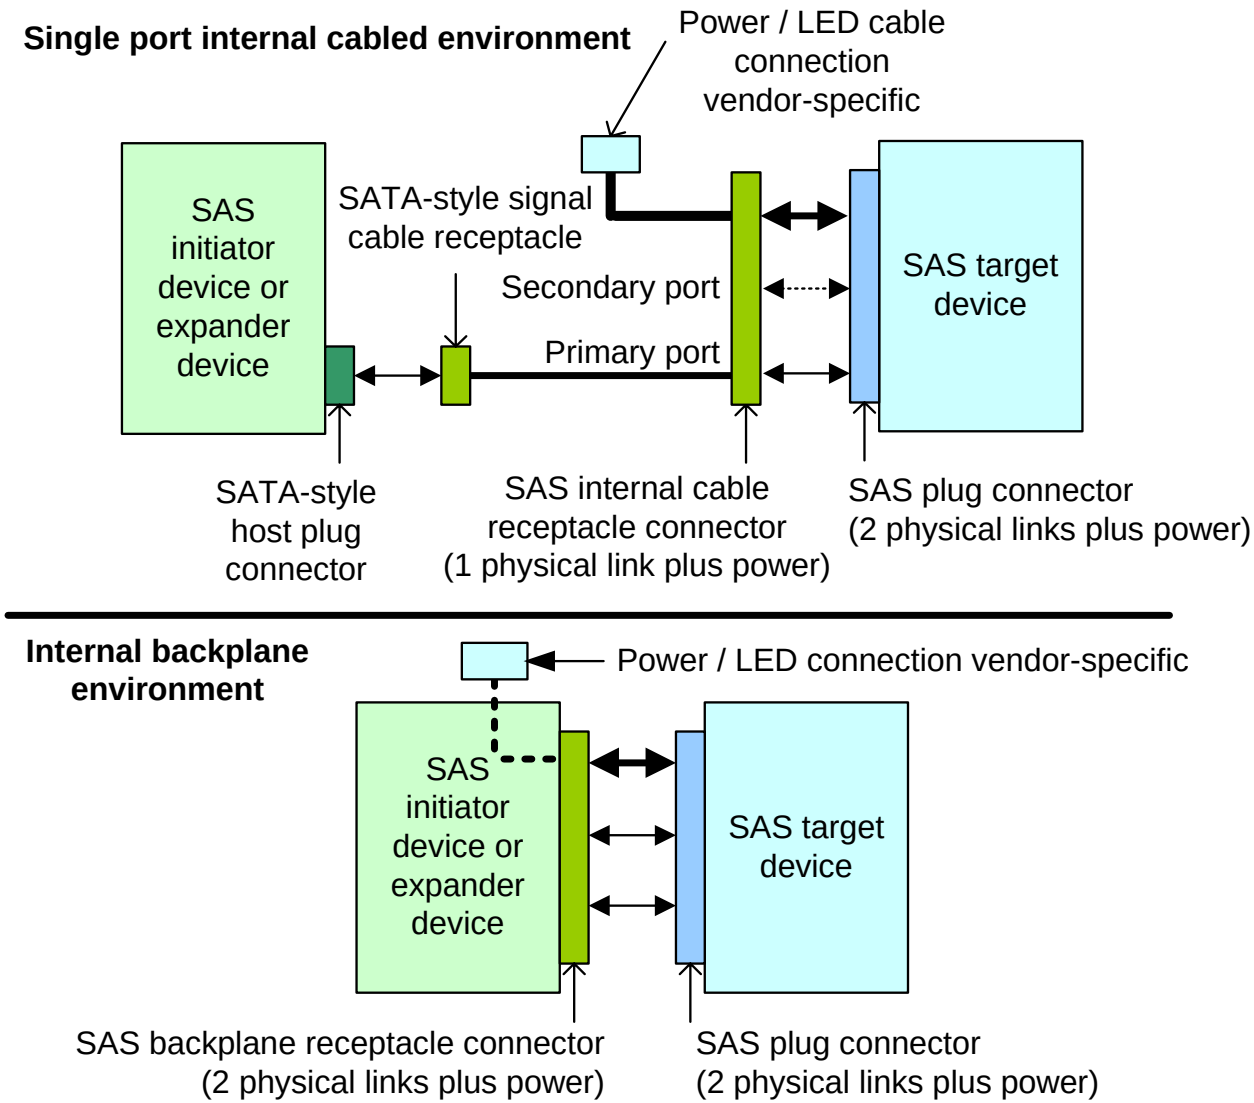

SAS plug and backplane receptacle connectors

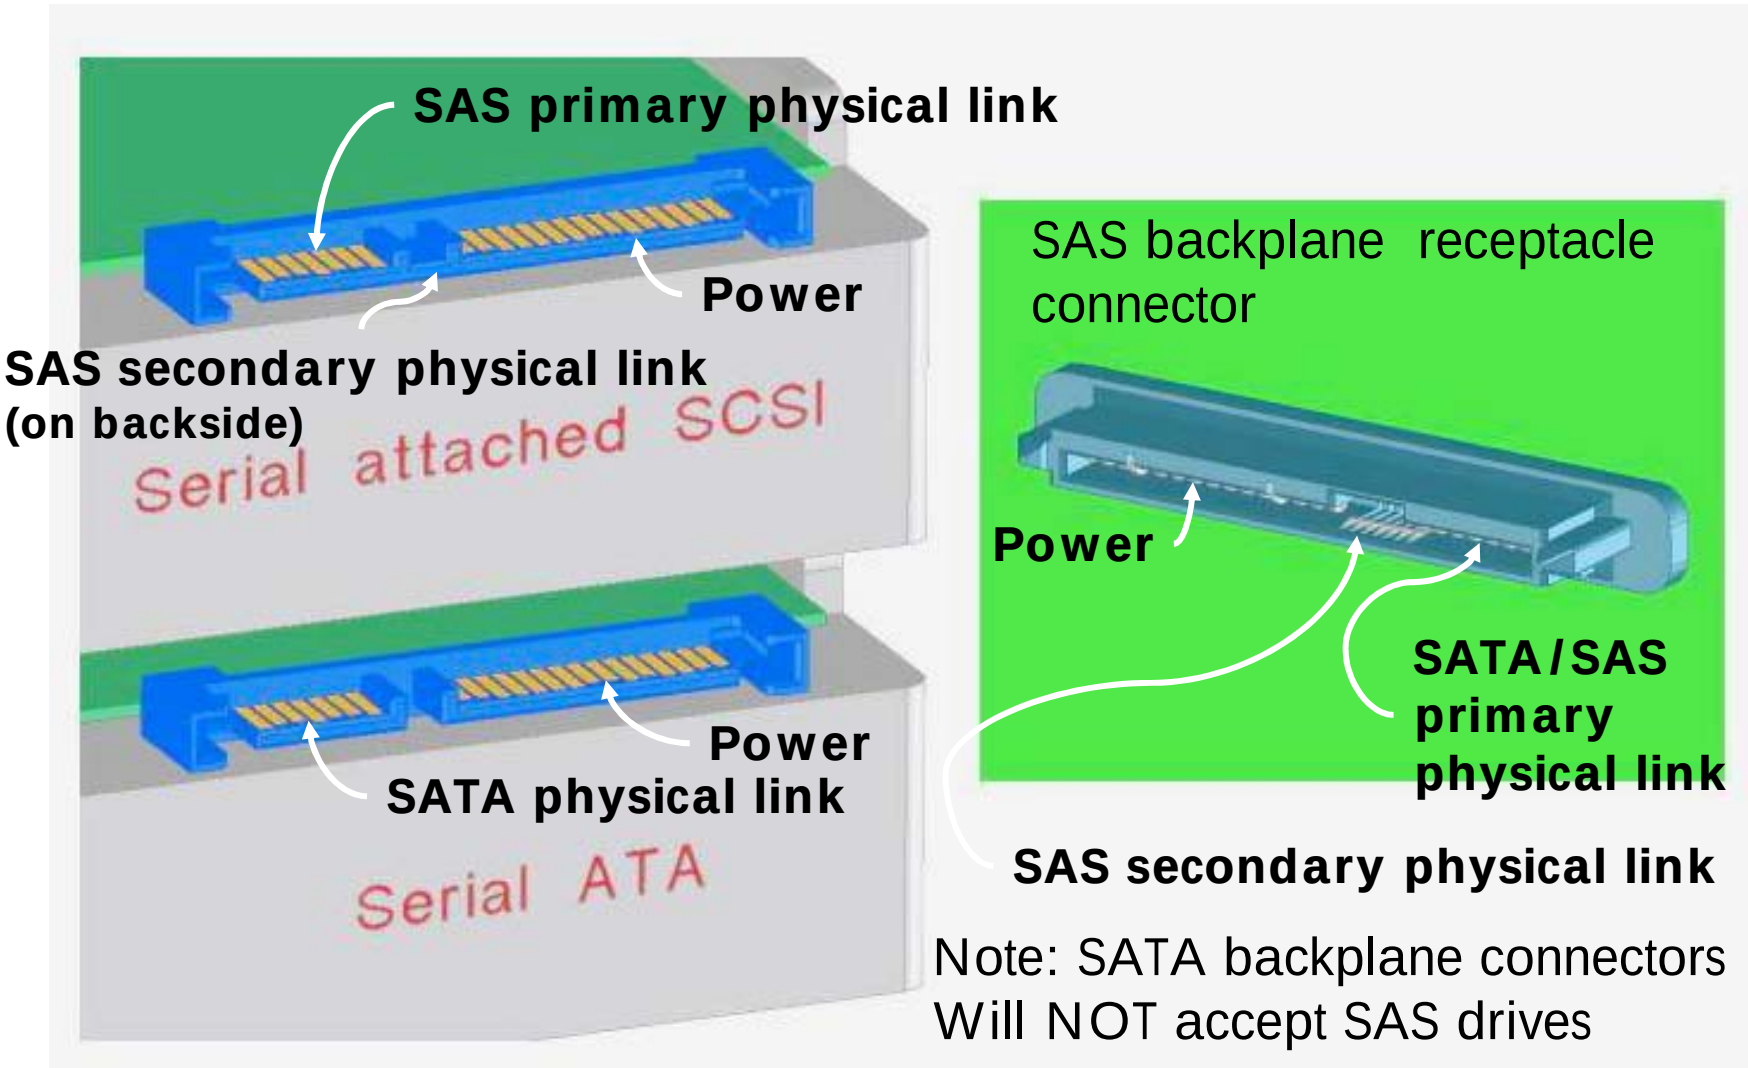

Note: SATA backplane connectors Will NOT accept SAS drives

Internal connector signal assignments

- Secondary physical link is optional
- The cable or backplane connects Tx on one side to Rx on the other side

Host connector	
Pin	Signal
S1	GROUND
S2	TP+
S3	TP-
S4	GROUND
S5	RP-
S6	RP+
S7	GROUND

Device connector		
Pin	Signal	Notes
S1	GROUND	Primary phy
S2	RP+	
S3	RP-	
S4	GROUND	
S5	TP-	
S6	TP+	
S7	GROUND	
S8	GROUND	Secondary phy (no-connects on SAS single-phy drives, SATA drives, and narrow cables)
S9	RS+	
S10	RS-	
S11	GROUND	
S12	TS-	
S13	TS+	
S14	GROUND	

Internal device connector power assignments

- Three voltages are provided
  - 3.3 V (new for serial interfaces)
  - 5 V
  - 12 V (for drive motors)
- All the pins of one voltage are tied together by the target device
- Precharge pins are longer
- Signal cables do not carry power
  - Separate power cable
  - Converters cables from old large power connector to the SATA hot plug capable power connector
- SFF discussing 12 V only drives

Pin	Signal
P1	V3.3
P2	V3.3
P3	V3.3, precharge
P4	GROUND
P5	GROUND
P6	GROUND
P7	V5, precharge
P8	V5
P9	V5
P10	GROUND
P11	READY LED
P12	GROUND
P13	V12, precharge
P14	V12
P15	V12

Internal wide connector

- Essentially 5 SATA connectors side by side
- Middle one used for sideband signals (e.g., I^2C or vendor-unique signals)
- Supports round or ribbon cable
- No power pins; grounds between pairs

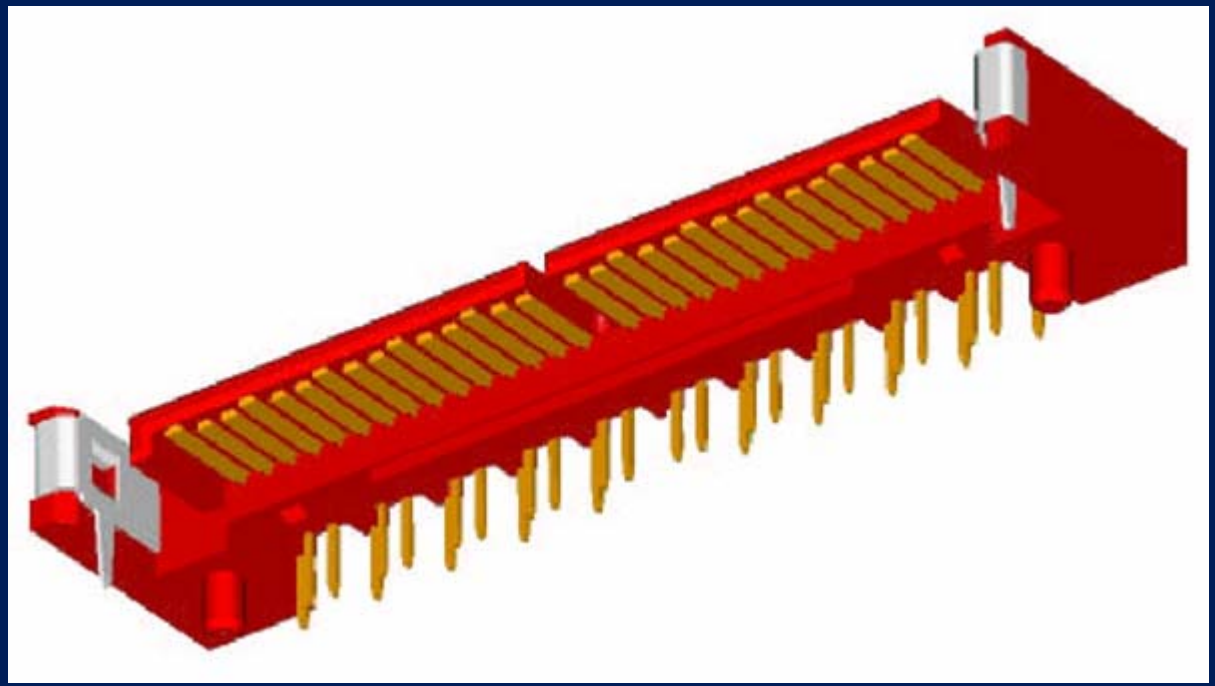

Pin	Signal
S1	Tx 0+
S2	Tx 0-
S3	Rx 0-
S4	Rx 0+
S5	Tx 1+
S6	Tx 1-
S7	Rx 1-
S8	Rx 1+
SB0	Sideband
SB1	Sideband
SB2	Sideband
SB3	Sideband
S9	Tx 2+
S10	Tx 2-
S11	Rx 2-
S12	Rx 2+
S13	Tx 3+
S14	Tx 3-
S15	Rx 3-
S16	Rx 3+

SAS cables and connectors - external

- SAS defines an external (box-to-box) environment using InfiniBand 4x connectors and cables

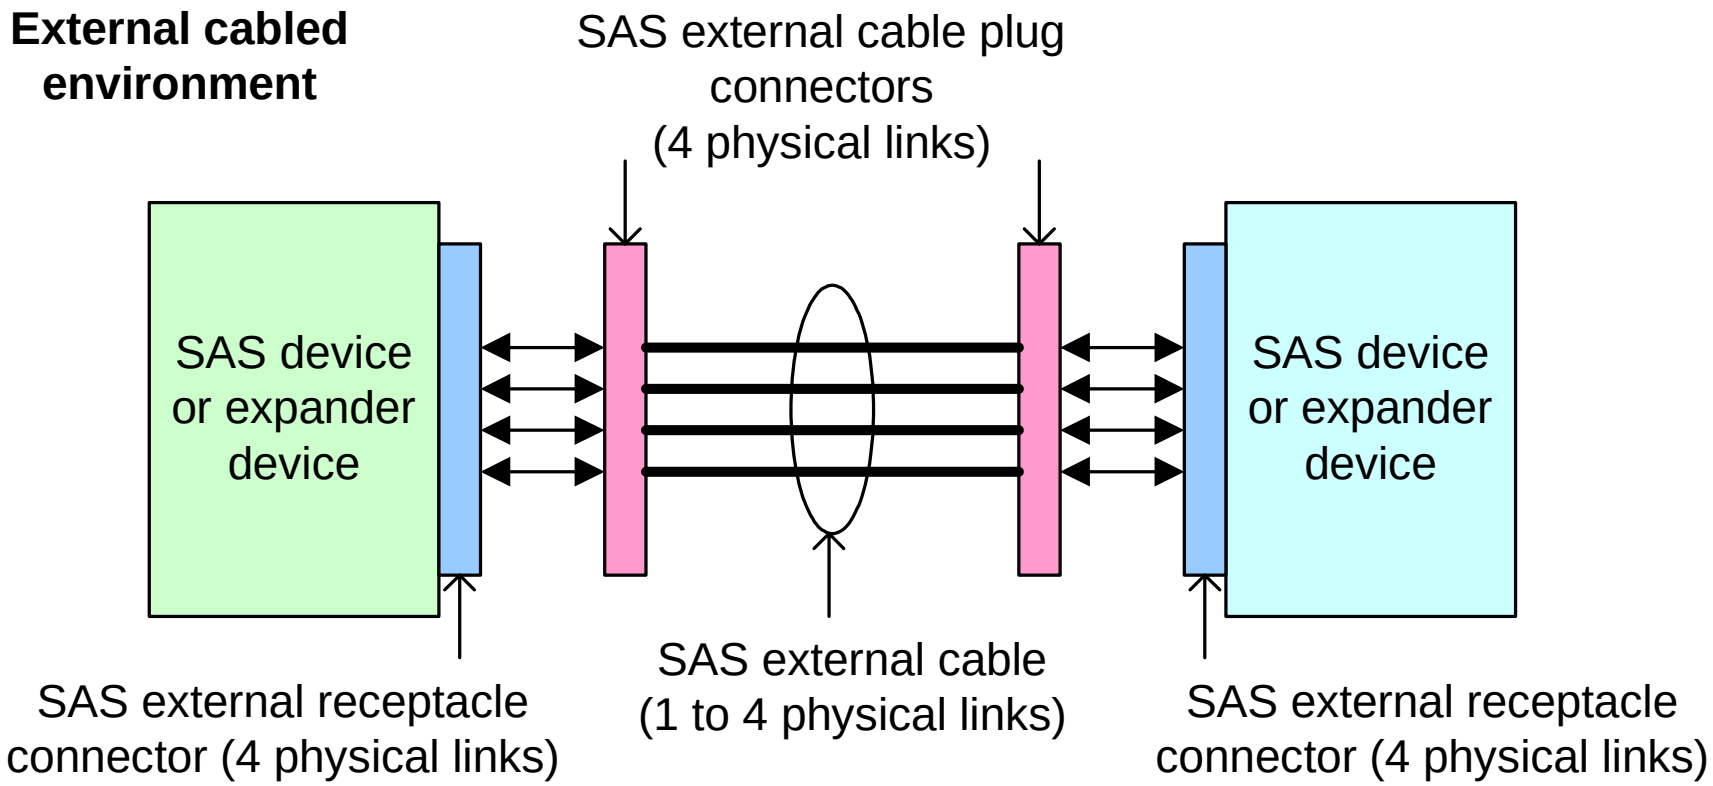

(SAS external cable connects the Tx signal pins to the Rx signal pins on each physical link)

SAS external receptacle connector

- (picture courtesy FCI / Berg)

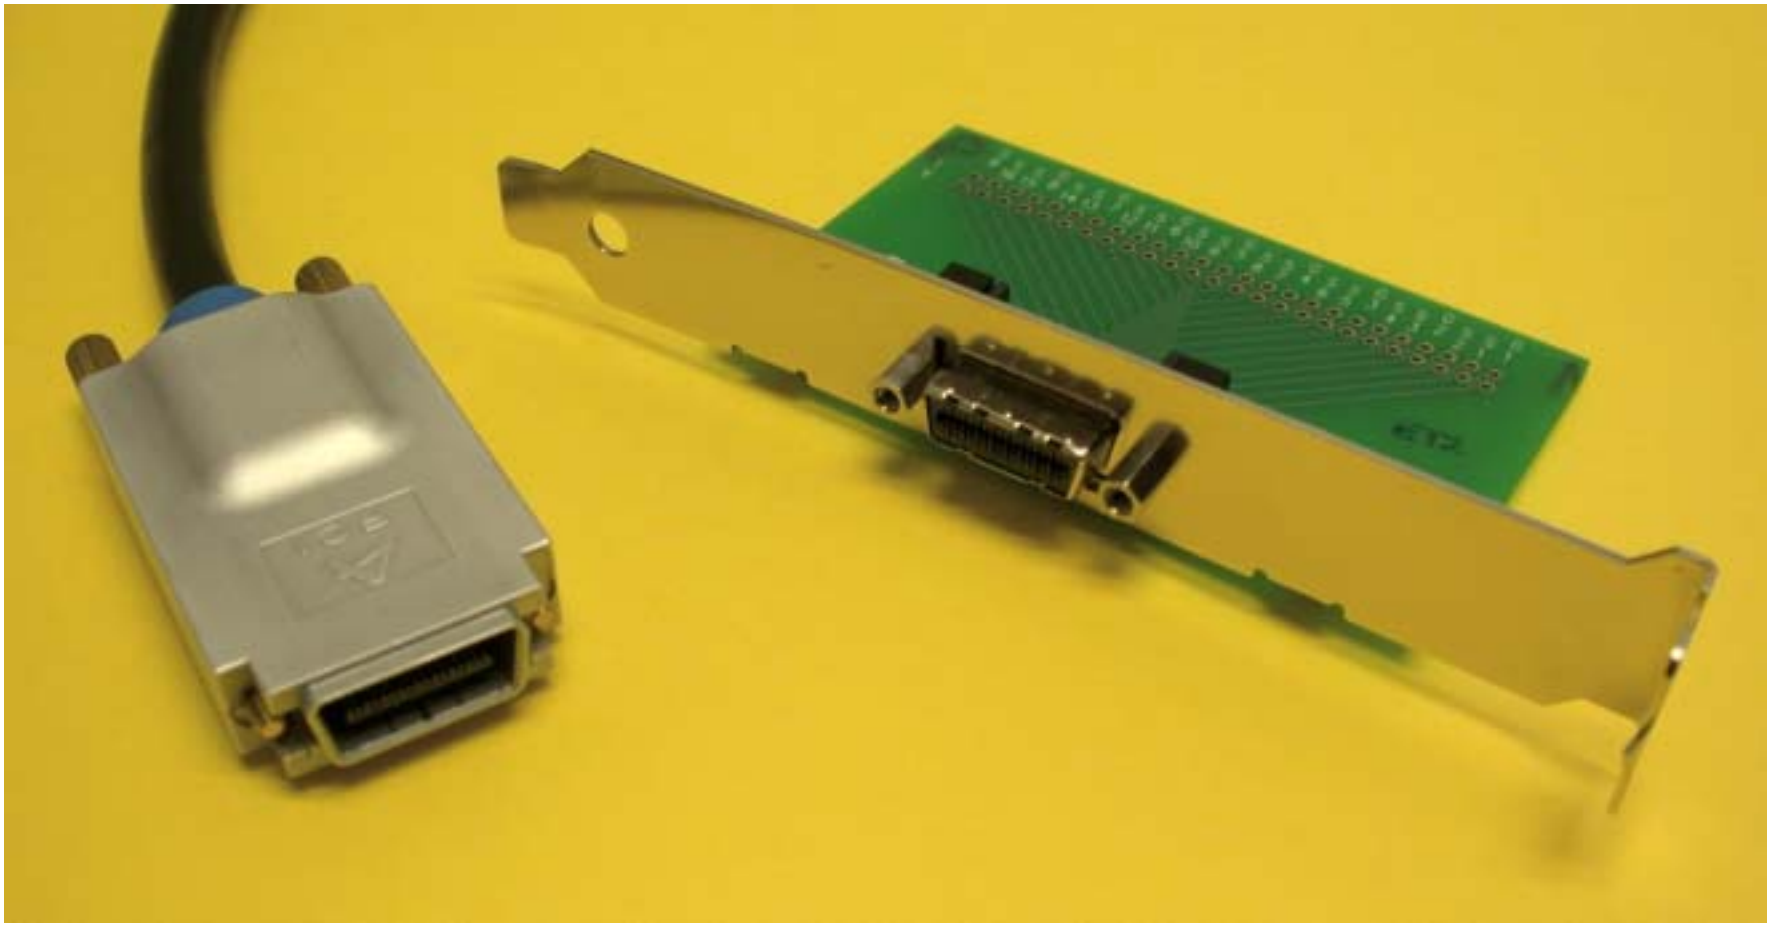

External connector signal assignments

- All the Rx signals are grouped together
- All the Tx signals are grouped together
- There are also 9 ground signals
- No external power signals
- No external READY LED signals
- Cable connects Tx on one side to Rx on the other side

Pin	Signal
S1	Rx 0+
S2	Rx 0-
S3	Rx 1+
S4	Rx 1-
S5	Rx 2+
S6	Rx 2-
S7	Rx 3+
S8	Rx 3-
S9	Tx 3-
S10	Tx 3+
S11	Tx 2-
S12	Tx 2+
S13	Tx 1-
S14	Tx 1+
S15	Tx 0-
S16	Tx 0+

Connector standards

- SFF defines SAS connectors and cables; SAS defines signal assignments

Standard	Description
SATA	SATA device plug and SATA signal host plug SATA host receptacle and SATA signal cable receptacle
SFF-8482	SAS plug (2x) SAS internal cable receptacle (1x or 2x) SAS backplane receptacle (2x)
SFF-8470	SAS external cable plug (4x) SAS external receptacle (4x)
SFF-8484	SAS internal wide plug and receptacle (4x)
SFF-8460	High-speed signaling backplane design guidelines
SFF-8223	2.5” form factor (connector location)
SFF-8323	3.5” form factor (connector location)
SFF-8523	5.25” form factor (connector location)

Physical layer - Electrical characteristics

Electrical characteristics

- SATA
  - Transmitters and cables are specified
  - Receivers are only indirectly (or imprecisely) specified
  - Both A.C. and D.C. coupling allowed
  - Spread spectrum clocking is optional
- SAS
  - Receiver tightly specified
    - Eye masks
    - 75 ps skew between Rx – and Rx +
  - Transmitter indirectly specified
    - Compliance channel approach tests transmitter under theoretical worst load or worse-than-theoretical actual load
  - A.C. coupling only (receivers must, transmitters may)
  - No SSC

Identical SAS/SATA characteristics

Characteristic	Value
Media impedance (differential)	100 ohm
Maximum AC coupling	12 nF
Maximum transmitter transients	± 1.2 V
Receiver sinusoidal AC common mode frequency tolerance range	2 to 200 MHz

- Based on same high-speed serial transceivers as other interfaces
  - Fibre Channel, Gigabit Ethernet, XAUI, InfiniBand, 1394b, PCI Express
- Many ASIC vendors' phys support all these interfaces

SATA Electrical specifications

Characteristic	Value at 1.5 Gbps
Transmitter differential voltage	400 to 600 mV(P-P)
Receiver differential voltage	325 to 600 mV(P-P)
Open circuit DC voltage level of each signal when AC coupling is used	0 to 2.0 V
Minimum OOB signal amplitude	50 to 100 mV(P-P)
Receiver sinusoidal AC common mode voltage tolerance	100 mV (P-P)

Characteristic	Value at 1.5 Gbps
Unit Interval (UI)	666.43 to 670.12 ps
Unit Interval during OOB signals	646.67 to 686.67 ps
Transmitter rise/fall time	0.15 to 0.41 UI
Transmitter differential skew	20 ps
Transmitter DC clock frequency variation	-350 to +350 ppm
Transmitter AC clock frequency variation Spread Spectrum Clocking - SSC	-5000 to 0 ppm

SAS compliance points

- Compliance points are where SAS places various requirements
- Most specifications are on receiver compliance points, not transmitters

Point	Description
IT, IR	Internal connectors
CT, CR	External connectors
XT, XR	Expanders (which may be attached to SATA drives)

General SAS interface characteristics

Characteristic	Value at 3 Gbps
Physical link rate	300 MB/sec
Unit interval (UI)	333.333 ps
Physical link rate tolerance at XR	+350/-5350 ppm
Physical link rate tolerance at IR and CR	± 100 ppm
Physical link rate tolerance at IT, CT, and XT	± 100 ppm
Maximum transients	± 1.2 V
Receiver AC common mode voltage tolerance	150 mV (P-P)

ppm = 0.01%

Transmitter Characteristic	Value at 1.5 Gbps	Value at 3 Gbps
Rise/fall time	67 to 273 ps	67 to 137 ps
Skew	20 ps	15 ps

Receiver eye mask

Characteristic	3 Gbps at IR, CR, XR
2 x Z2	1600 mV (P-P)
2 x Z1	275 mV (P-P)
X1	0.275 UI
X2	0.50 UI

Characteristic	1.5 Gbps at IR, XR
2 x Z2	1600 mV (P-P)
2 x Z1	325 mV (P-P)
X1	0.275 UI
X2	0.50 UI

Characteristic	SATA 1.5 Gbps
2 x Z2	600 mV (P-P)
2 x Z1	325 mV (P-P)

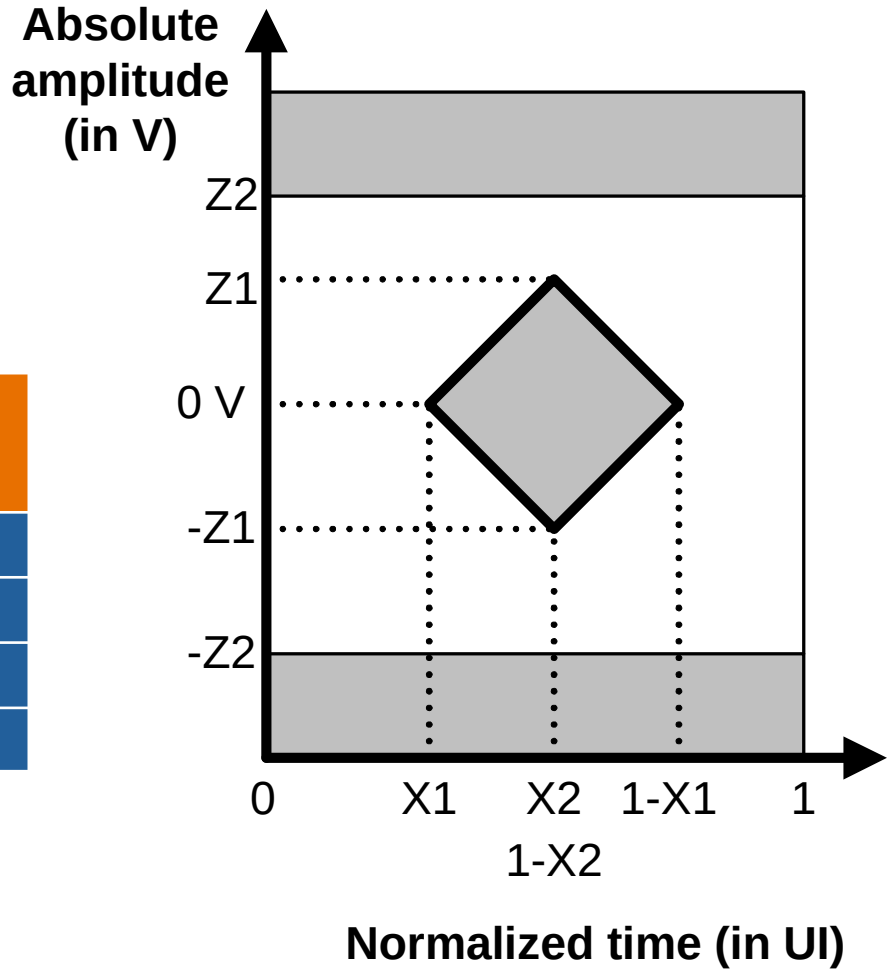

Jitter types

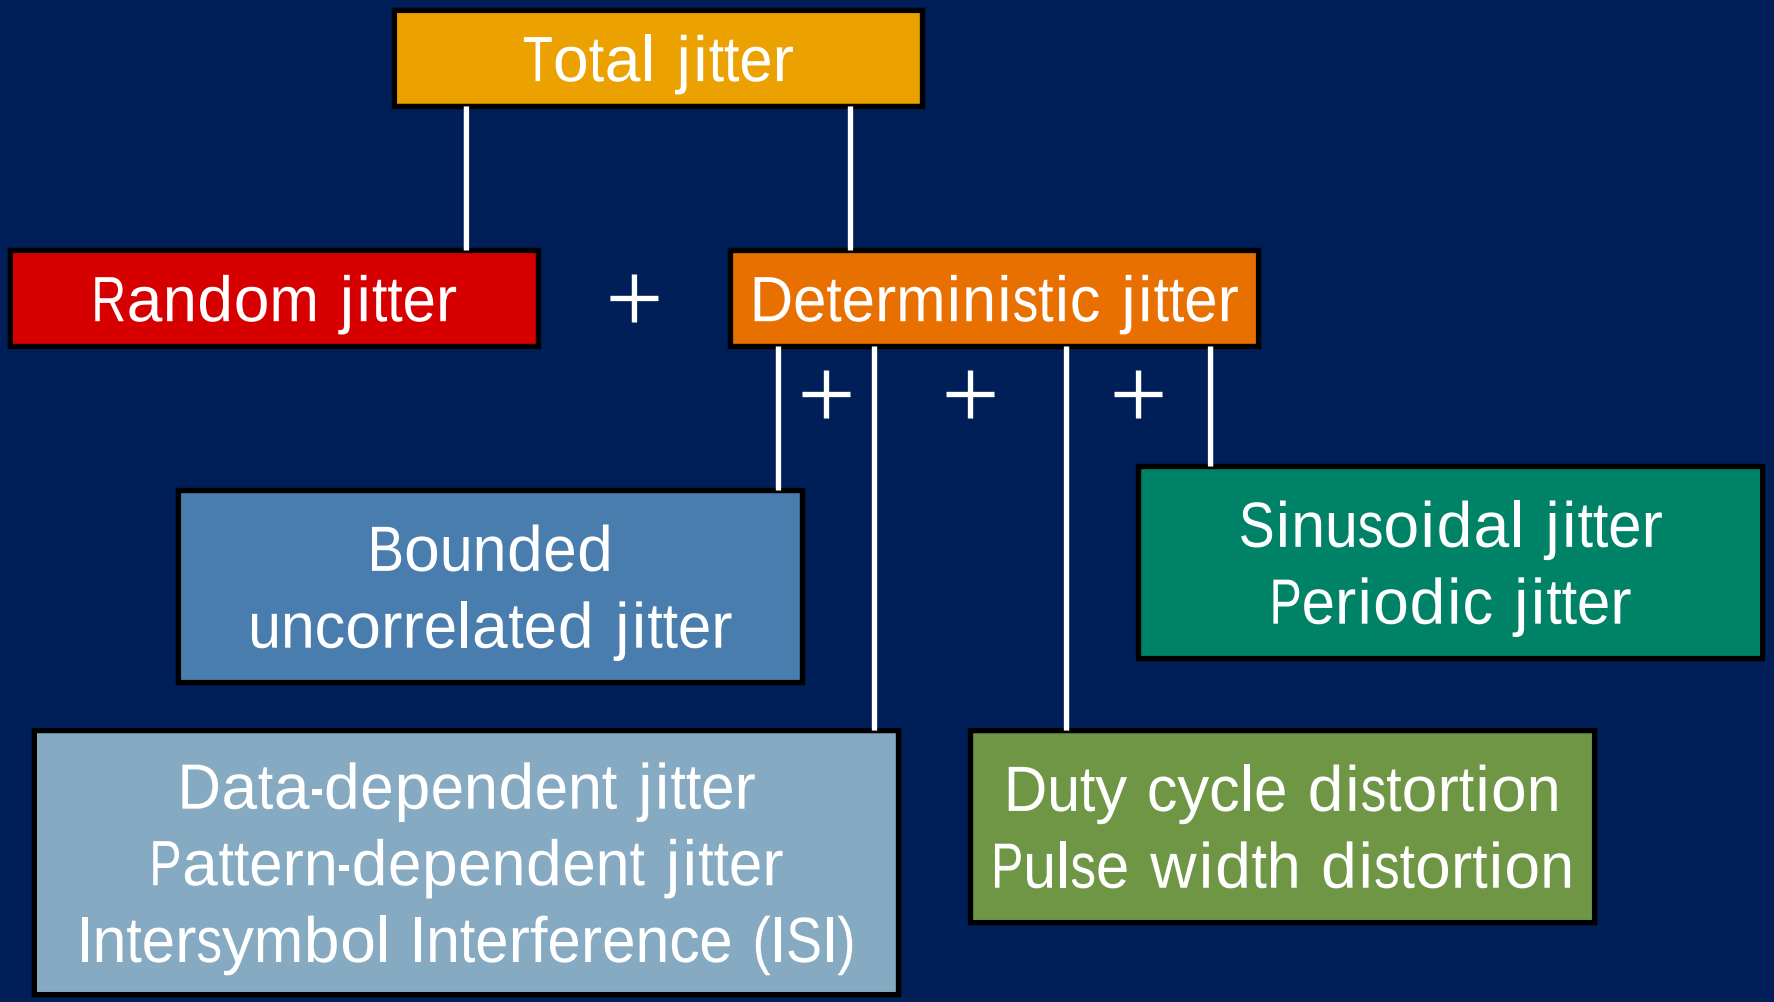

Jitter definitions

Type	Definition
Total jitter	Deviation of an event from its ideal time. Total jitter = Deterministic jitter + random jitter.
Deterministic jitter	Jitter component bounded in amplitude. Includes duty-cycle distortion, data-dependent, sinusoidal, and bounded uncorrelated jitter.
Duty-cycle distortion Pulse-width distortion	Difference between the duty cycle (bit period) of a 1 bit and a 0 bit.
Data-dependent jitter Pattern-dependent jitter Intersymbol interference (ISI)	Jitter that occurs on non-clocklike waveforms when the datapath's bandwidth limitations distort the waveform edge locations.
Sinusoidal jitter Periodic jitter	Jitter caused by signal crosstalk of signals in the system related to the data pattern.
Bounded uncorrelated jitter	Jitter that does not correlate to the data pattern but is bounded.
Random jitter	Jitter that is unbounded in amplitude and follows a Gaussian probability distribution. Mean is 0; standard deviation expressed in rms; infinite tails.

Jitter requirements

- Receivers must tolerate specified amounts of jitter using the CJTPAT test pattern

Characteristic	Value at 3 Gbps	Notes
Sinusoidal	0.10 UI	Caused by crosstalk of related signals
Deterministic *	0.35 UI	Duty-cycle distortion, data-dependent jitter (intersymbol interference), sinusoidal jitter, and uncorrelated jitter (unrelated signals)
Total allowed at receiver compliance point *	0.55 UI	Deterministic + random
Total that receivers shall tolerate *	0.65 UI	Deterministic + random

- Considered after applying a high-pass function that attenuates low frequency jitter

Transmitter characteristics

- TCTF = Transmitter compliance transfer function
- Mathematical model of a system through which a transmitter shall be capable of delivering signals that meet the receiver eye mask
- Two compliance techniques
  - Zero-length test load
    - Measure signal through this test load and mathematically apply the TCTF transfer function
  - Transmitter compliance transfer function test load
    - Measure signal through this a test load worse than any real system

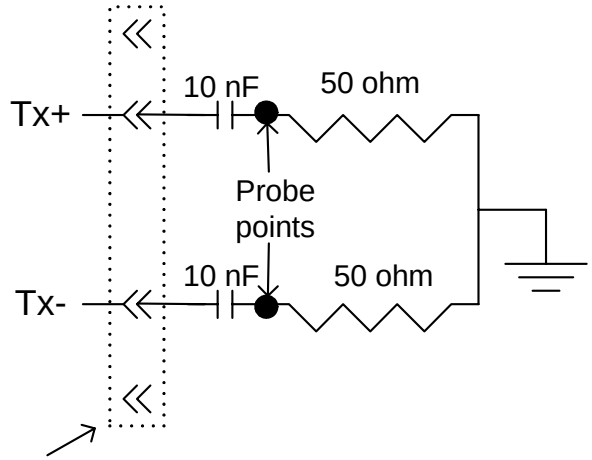

Intra-enclosure uses the SAS internal connector.
Inter-enclosure uses the SAS external connector.

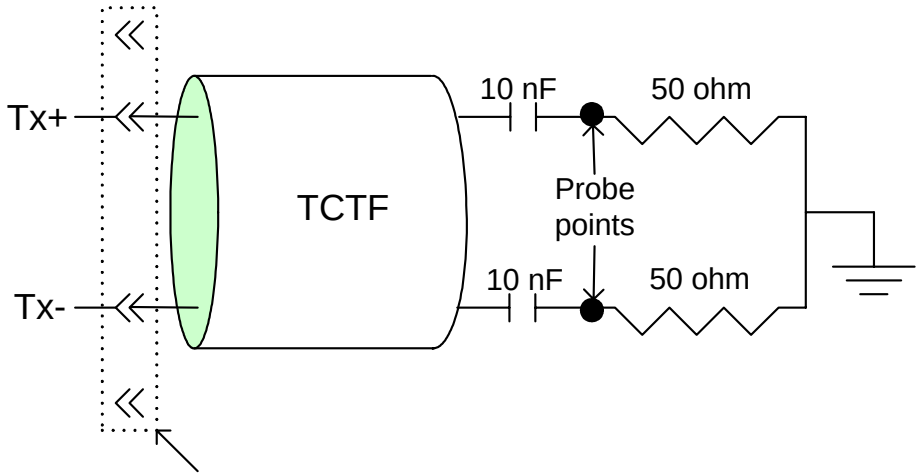

Intra-enclosure uses the SAS internal connector.
Inter-enclosure uses the SAS external connector.

READY LED signal electrical characteristics

- SAS target devices can drive their own activity LEDs
- SATA devices do not have this output
- Open collector or open drain circuit (drive low, tristate high)

State	Test condition	Requirement
Negated	$0 < V_{OH} < 3.6$ V	-100 uA $< I_{OH} < 100$ uA
Asserted	$I_{OL} = 15$ mA	$0 < V_{OL} < 0.225$ V

Effects of interconnect on signals

- The interconnect introduce more loss at high-frequencies than low frequencies
  - View as the interconnect takes time to fully “charge up”
- Conductor and dielectric losses
- Attenuation and phase shift

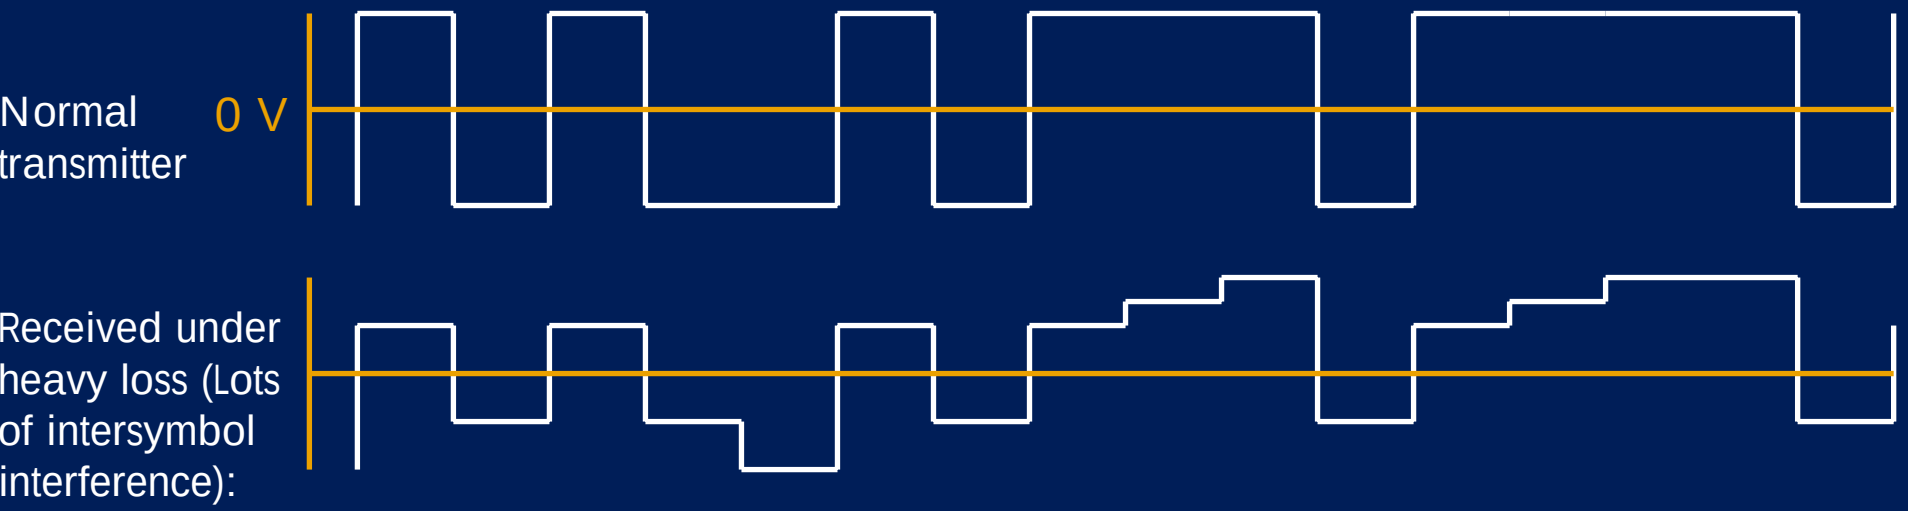

Interconnect loss example – short cable

- After going through a 1 m cable

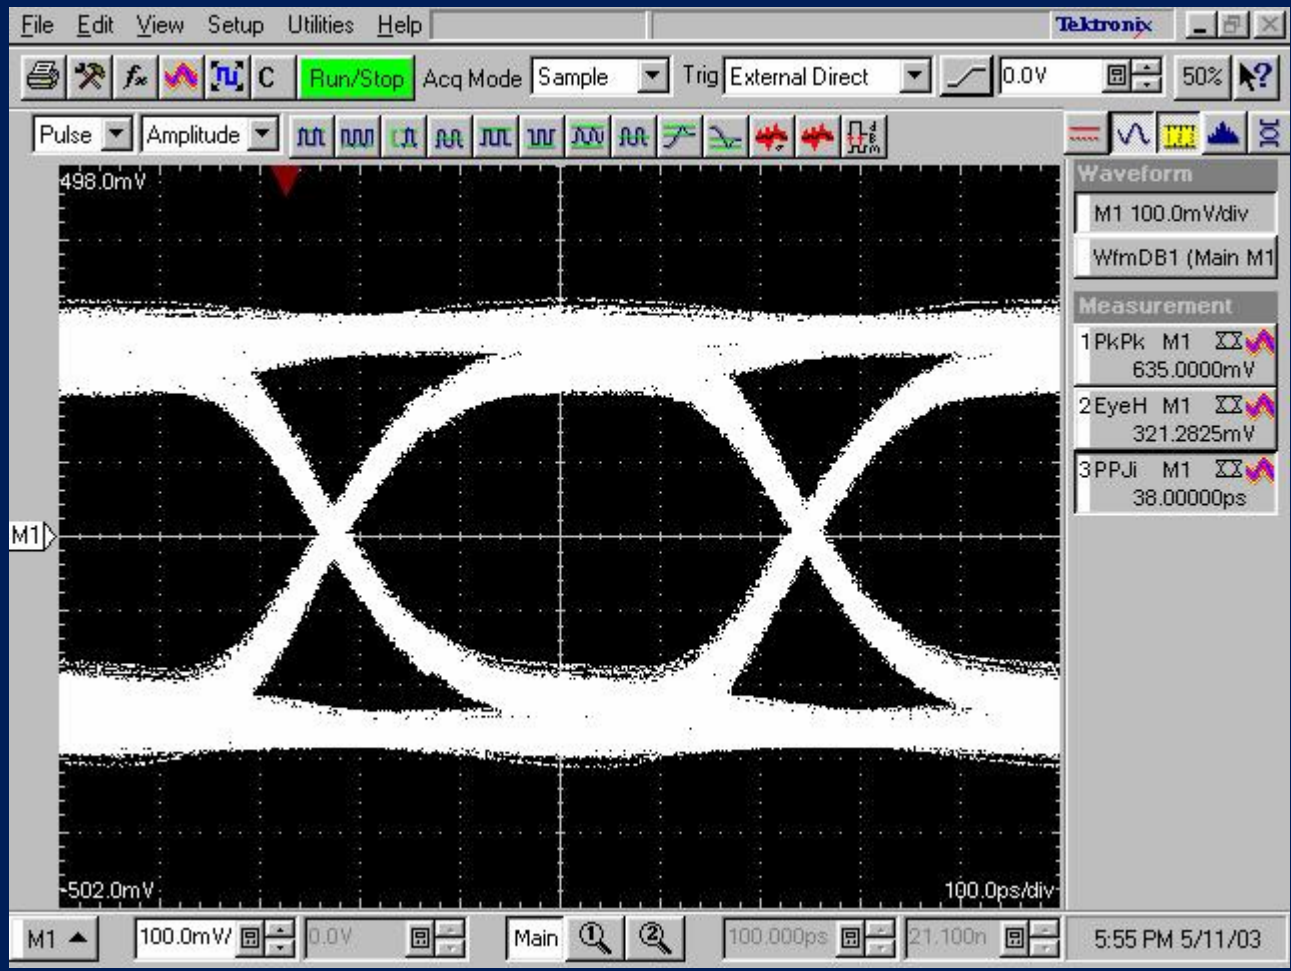

Interconnect loss example – long cable

- After going through a 10 m cable

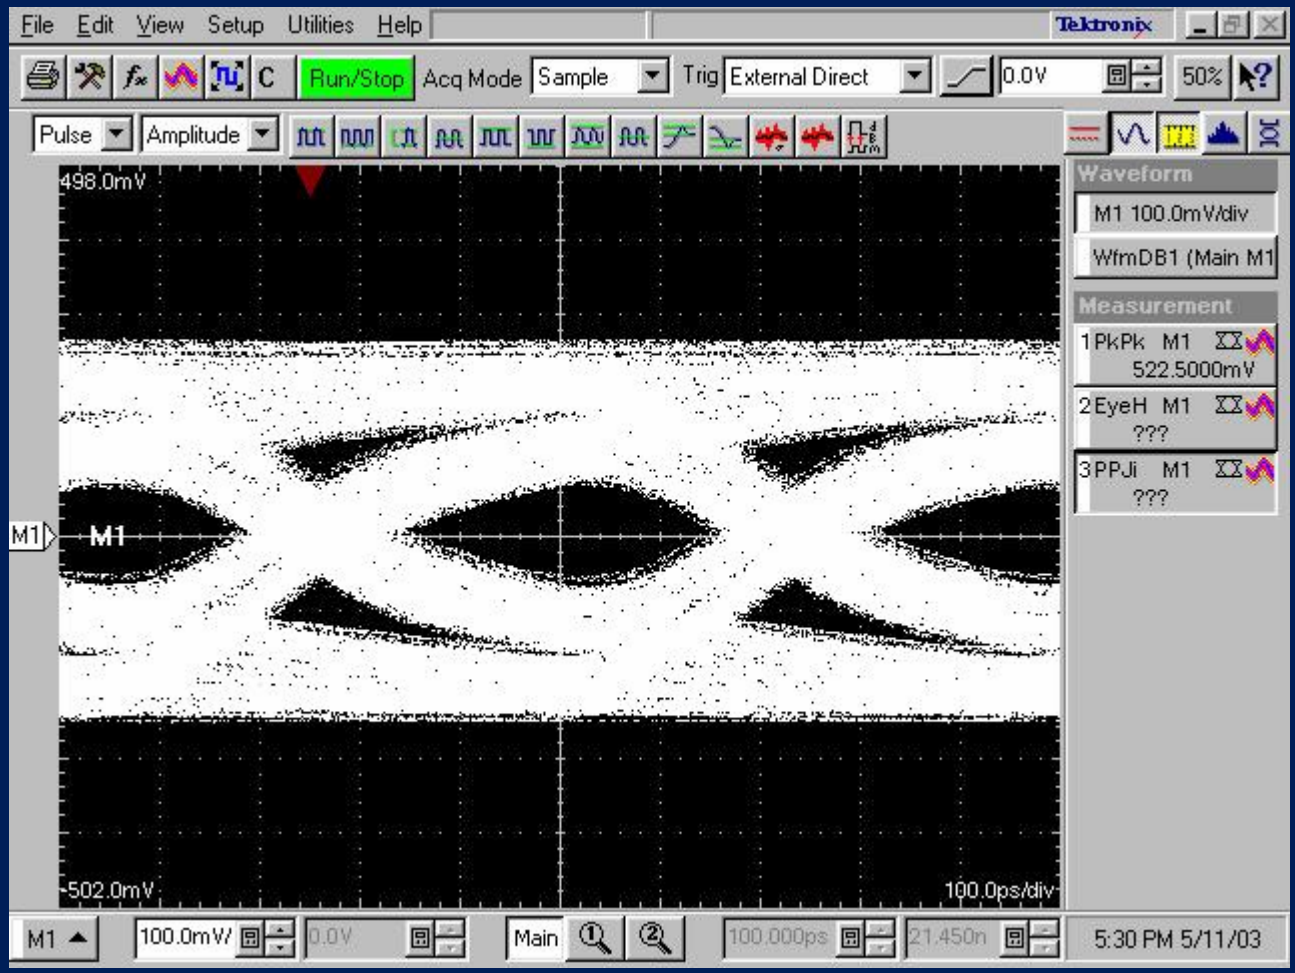

Countering interconnect loss

- Transmitter preemphasis
  - Change the signal going into the interconnect to counter the expected interconnect loss
- Receiver equalization
  - Change interpretation of signal received realizing interconnect has modified it

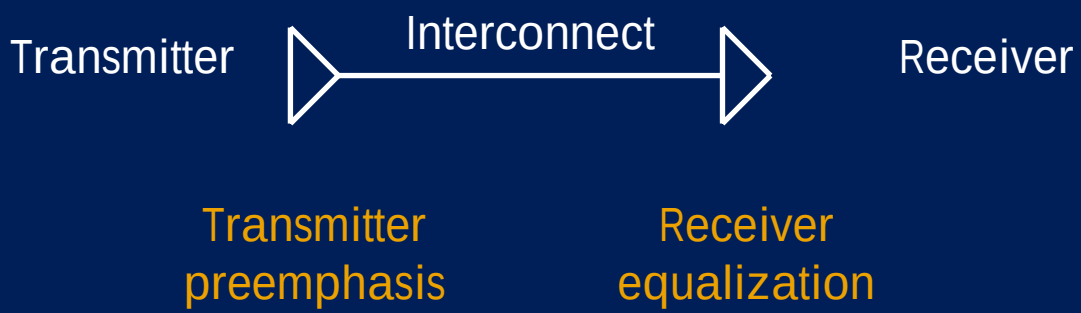

Transmitter preemphasis

- Amplify higher frequencies to counter cable loss
- Drive signals full strength after transition
- Drive signals weaker (up to 40%) when consecutive bits stay the same
- Finite Impulse Response (FIR) or Infinite Impulse Response (IIR) filters
  - FIR = based on historical input
  - IIR = based on historical input and output
  - More complex filters better counter the cable loss effects
- Optional (barely mentioned in either standard)

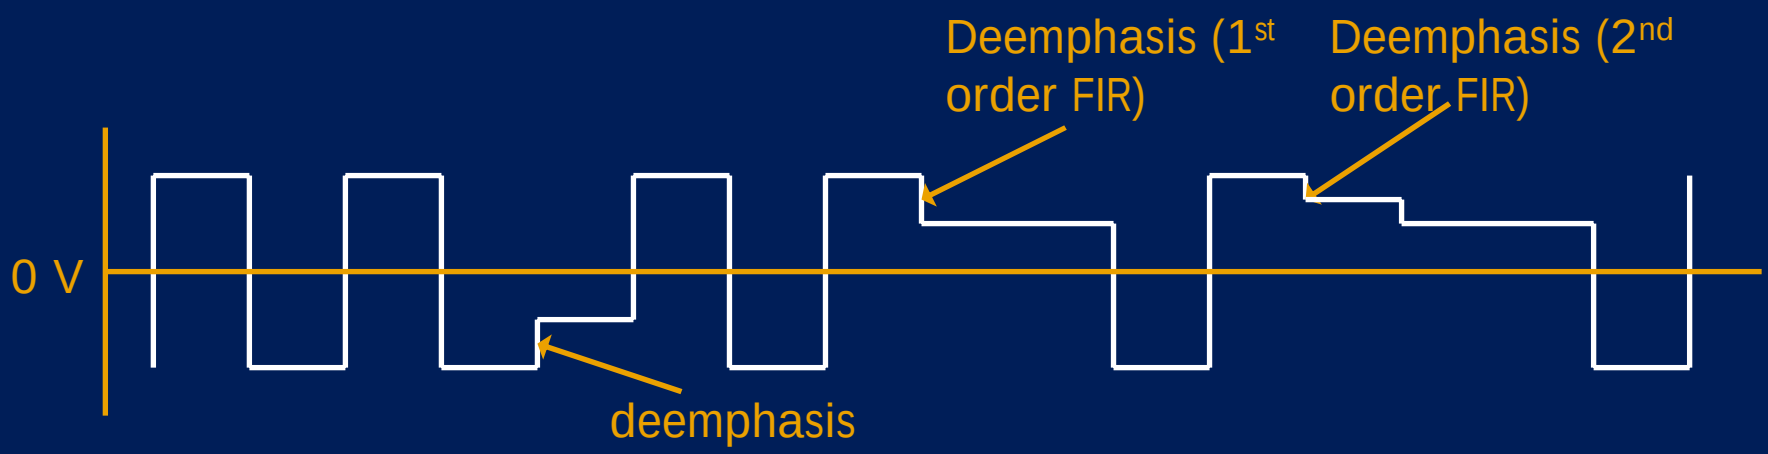

Transmitter preemphasis example

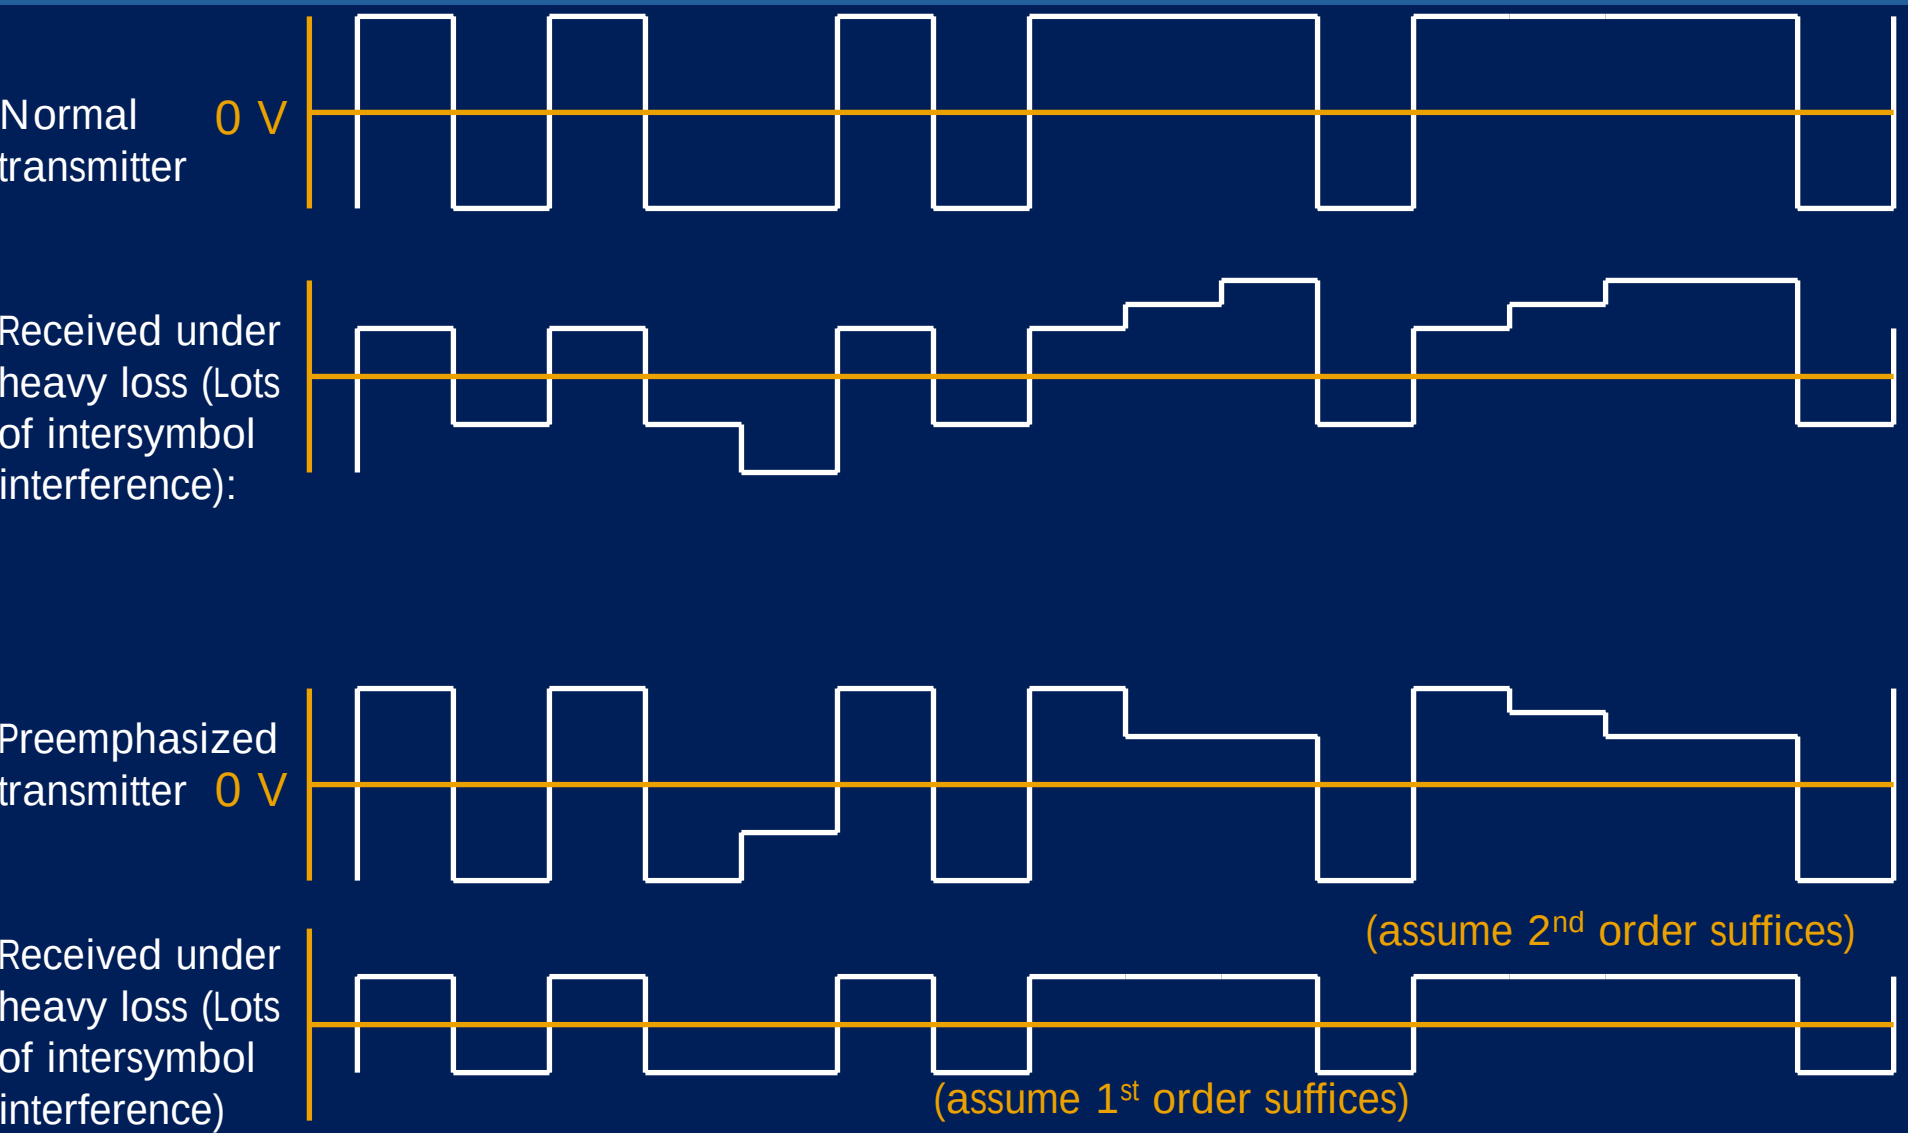

Preemphasis example – short cable

- Preemphasis with a 1 m cable

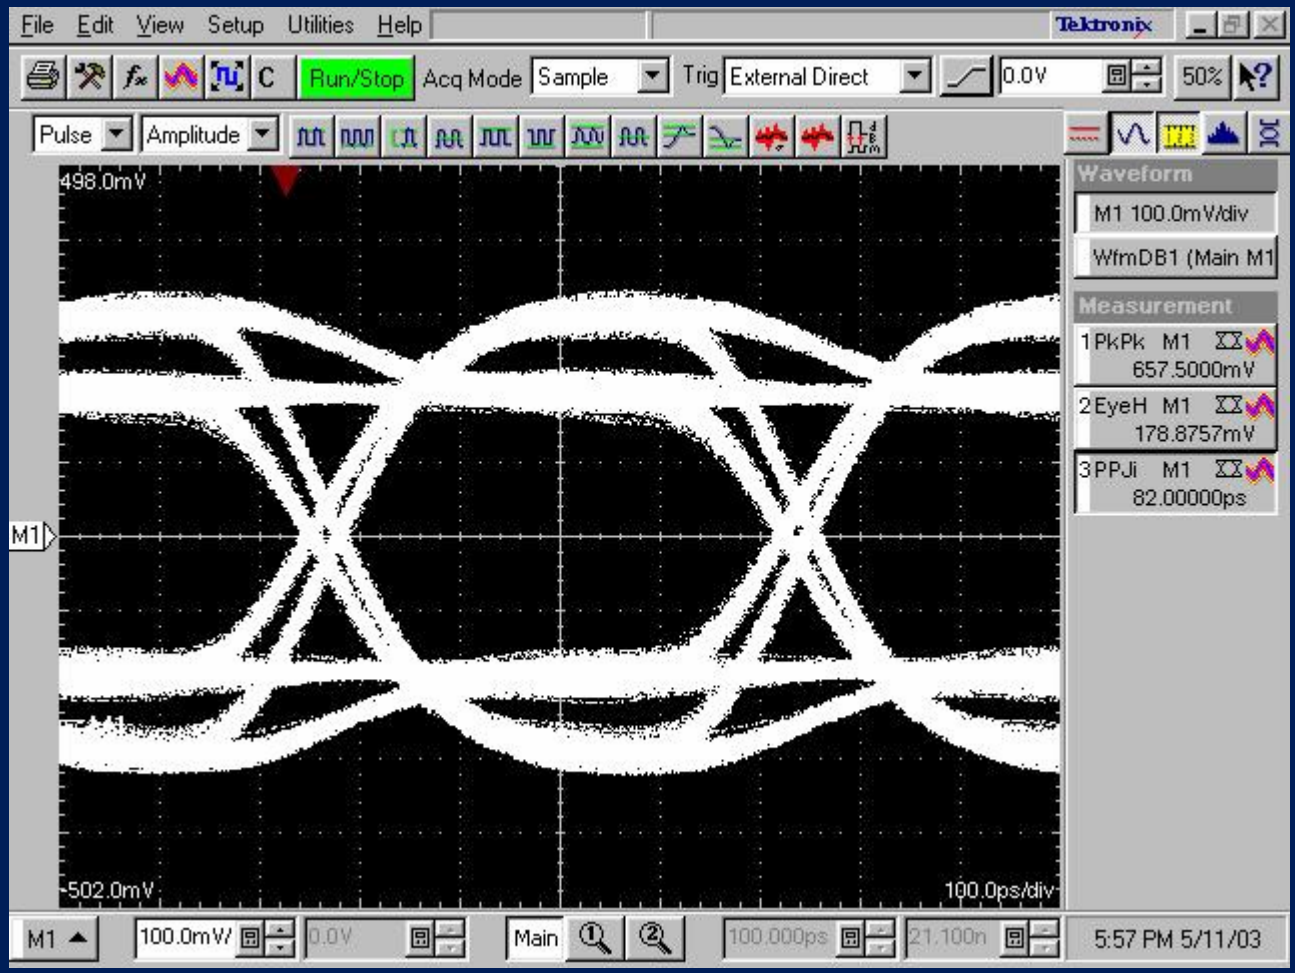

Preemphasis example – long cable

- Preemphasis with a 10 m cable

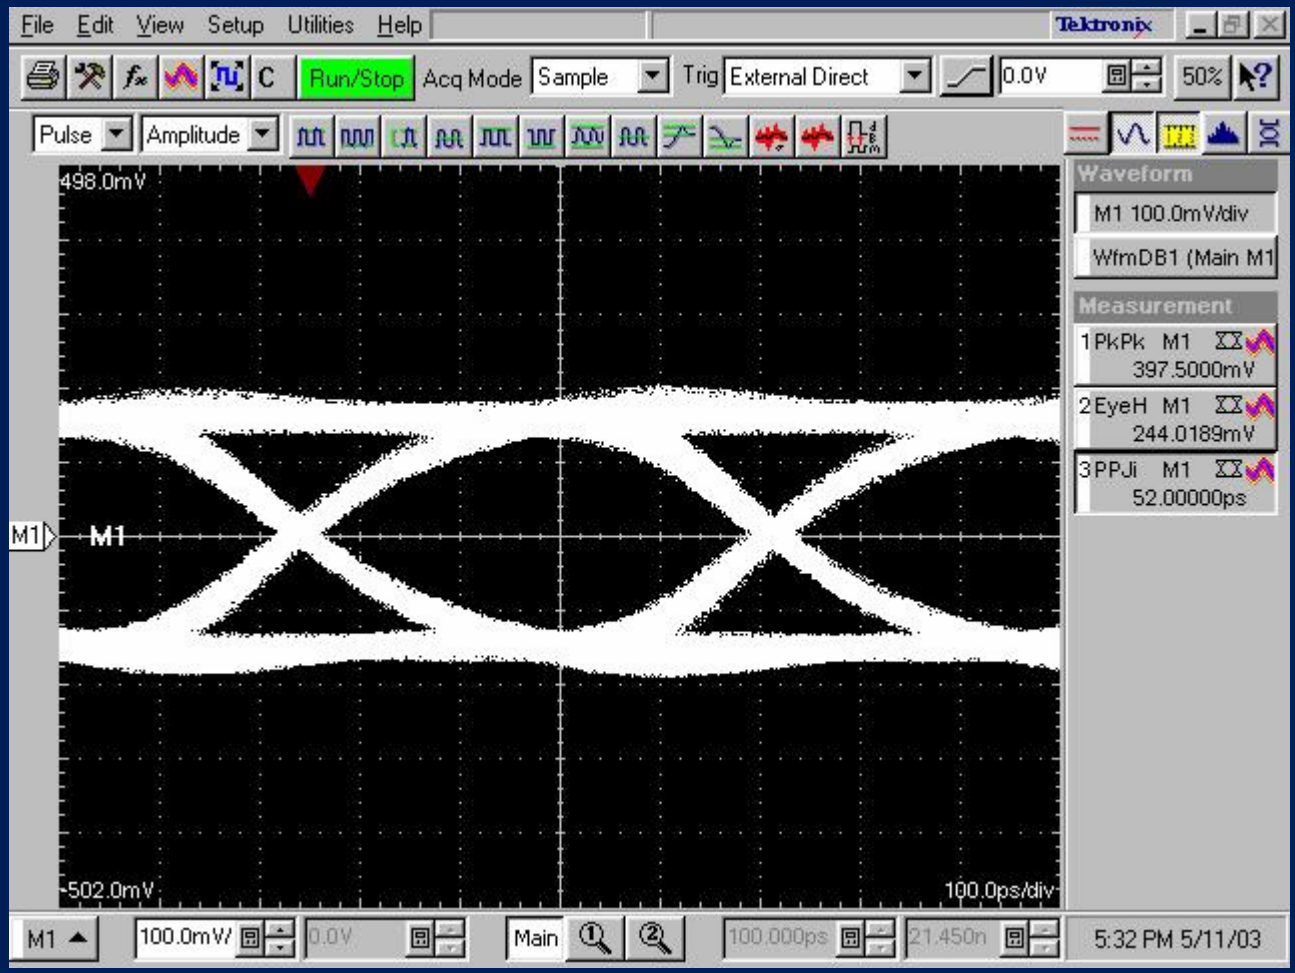

Receiver equalization

- Amplify higher frequencies to counter cable loss
  - Amplify (boost) input frequencies near 1.5/3.0 GHz
- Filter out higher and lower frequencies
- Can even recover signals when no eye is present at the receiver
- Pros – can automatically adapt to environment
- Cons – amplifies noise, cannot observe receiver behavior
- Optional (barely mentioned in either standard)

Receiver equalization example

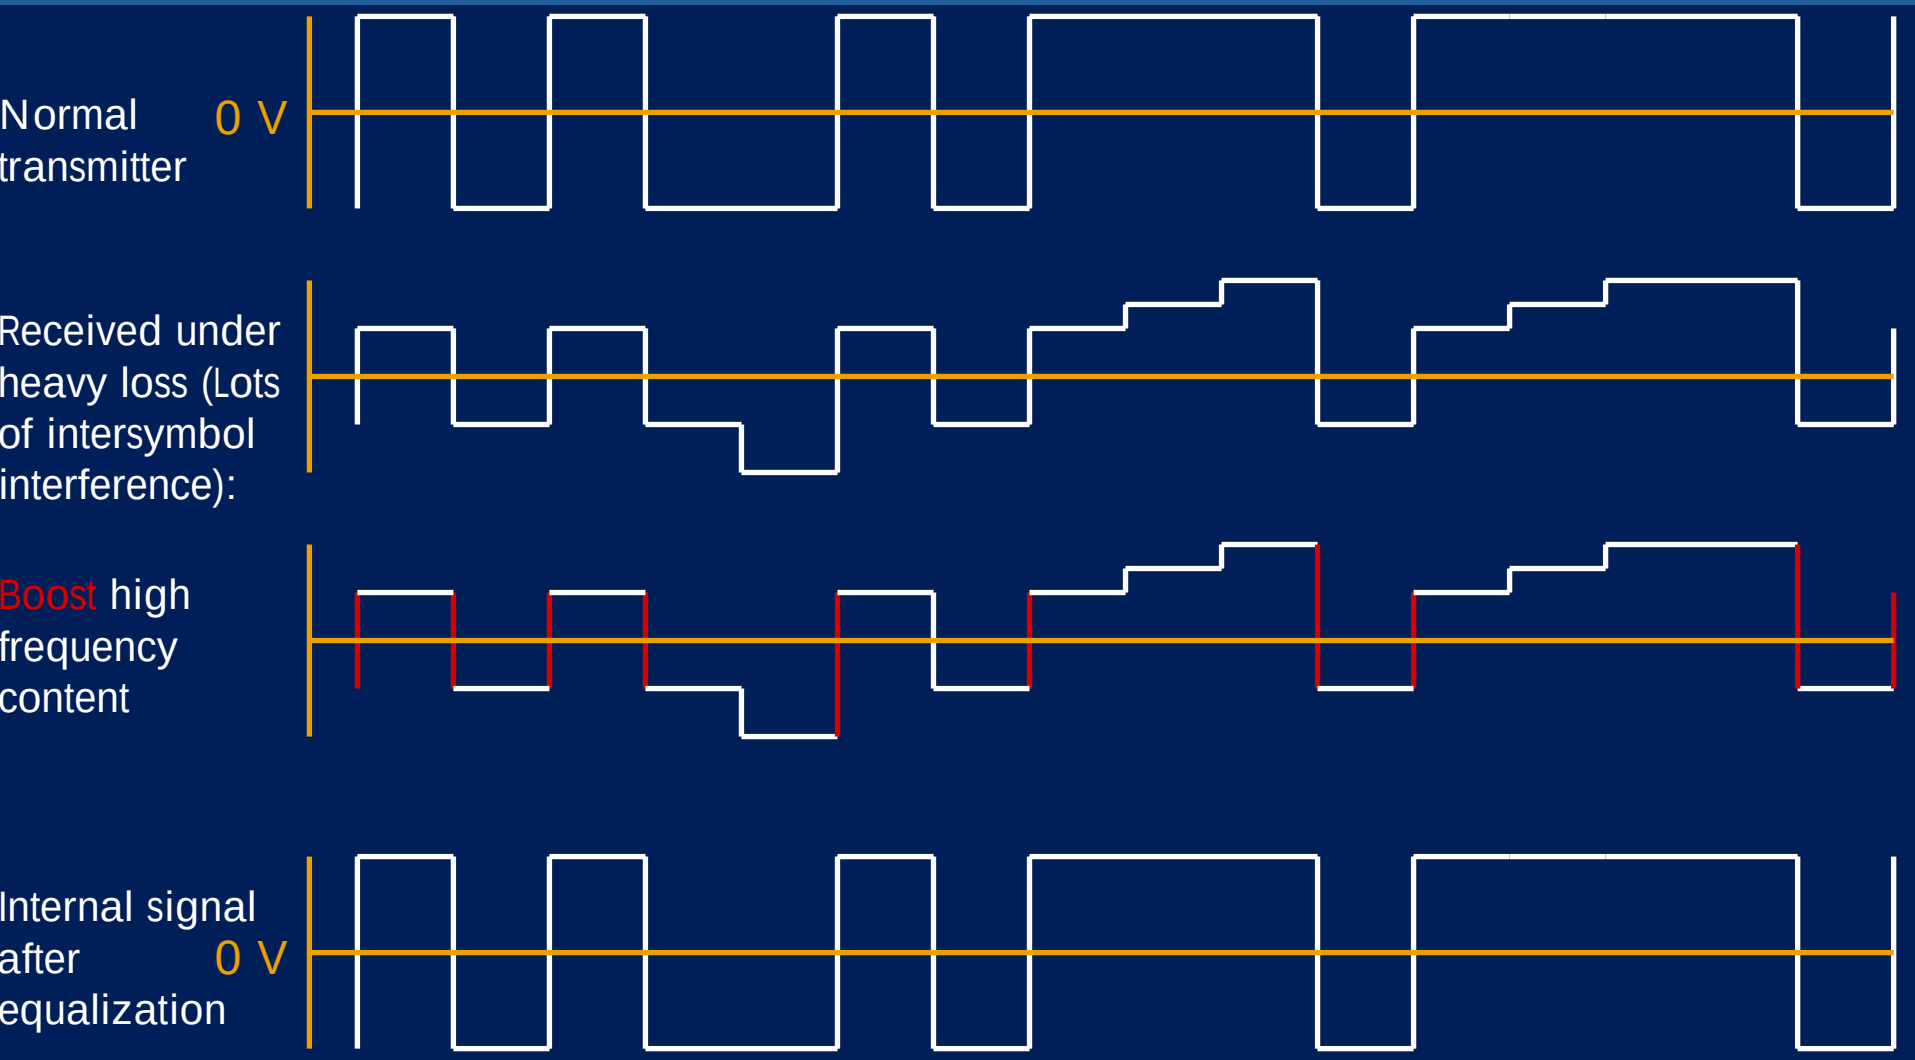

Additional references

- EDN magazine articles on http://www.edn.com
  - 6/10/1999: Pre-emphasis improves RS-485 communications
  - 5/1/2003: Eyeing jitter: shaking out why signals shake
  - 10/17/2002: An eye-opening look at jitter
  - 11/28/2002: Testing gigabit serial buses: First, get physical
  - 10/31/2002: Protecting high-speed buses at 1 Gbps and beyond
  - 3/6/2003: Receiver equalization increases link distance without adding EMI
- Howard Johnson's articles archived at http://www.signalintegrity.com

Wrap up

Serial Attached SCSI tutorials

- General overview (~2 hours)
- Detailed multi-part tutorial (~3 days to present):
  - Architecture
  - Physical layer
  - Phy layer
  - Link layer
    - Part 1) Primitives, address frames, connections
    - Part 2) Arbitration fairness, deadlocks and livelocks, rate matching, SSP, STP, and SMP frame transmission
  - Upper layers
    - Part 1) SCSI application and SSP transport layers
    - Part 2) ATA application and STP/SATA transport layers
    - Part 3) Management application and SMP transport layers, plus port layer
  - SAS SSP comparison with Fibre Channel FCP

Key SCSI standards

- Working drafts of SCSI standards are available on http://www.t10.org
- Published through http://www.incits.org
  - Serial Attached SCSI
  - SCSI Architecture Model – 3 (SAM-3)
  - SCSI Primary Commands – 3 (SPC-3)
  - SCSI Block Commands – 2 (SBC-2)
  - SCSI Stream Commands – 2 (SSC-2)
  - SCSI Enclosure Services – 2 (SES-2)
- SAS connector specifications are available on http://www.sffcommittee.org
  - SFF 8482 (internal backplane/drive)
  - SFF 8470 (external 4-wide)
  - SFF 8223, 8224, 8225 (2.5”, 3.5”, 5.25” form factors)
  - SFF 8484 (internal 4-wide)

Key ATA standards

- Working drafts of ATA standards are available on http://www.t13.org
  - Serial ATA 1.0a (output of private WG)
  - ATA/ATAPI-7 Volume 1 (architecture and commands)
  - ATA/ATAPI-7 Volume 3 (Serial ATA standard)
- Serial ATA II specifications are available on http://www.t10.org and http://www.serialata.org
  - Serial ATA II: Extensions to Serial ATA 1.0
  - Serial ATA II: Port Multiplier
  - Serial ATA II: Port Selector
  - Serial ATA II: Cables and Connectors Volume 1

For more information

- International Committee for Information Technology Standards
  - http://www.incits.org
- T10 (SCSI standards)
  - http://www.t10.org
  - Latest SAS working draft
  - T10 reflector for developers
- T13 (ATA standards)
  - http://www.t13.org
  - T13 reflector for developers
- T11 (Fibre Channel standards)
  - http://www.t11.org
- SFF (connectors)
  - http://www.sffcommittee.org
- SCSI Trade Association
  - http://www.scsita.org
- Serial ATA Working Group
  - http://www.serialata.org
- SNIA (Storage Networking Industry Association)
  - http://www.snia.org
- Industry news
  - http://www.infostor.com
  - http://www.byteandswitch.com
  - http://www.wwpi.com
  - http://searchstorage.com
- Training
  - http://www.knowledgetek.com


hp
invent