



# Serial Attached SCSI

## General overview



by Rob Elliott

HP Industry Standard Servers

Server Storage Advanced Technology

[elliott@hp.com](mailto:elliott@hp.com) <http://www.hp.com>

30 September 2003

- 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





# SCSI – Small Computer System Interface

# History of parallel SCSI

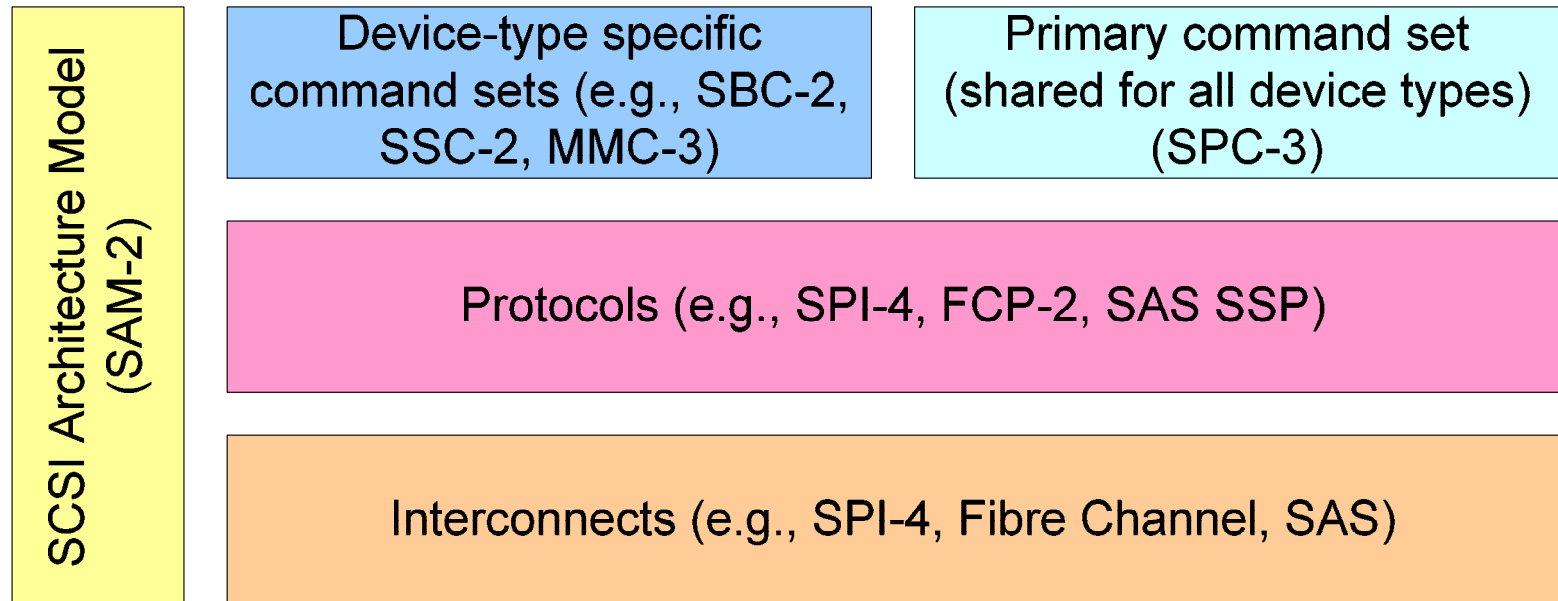


Interconnect	Standard	Year	Speed	Key features
SASI		1979		Shugart Associates
SCSI-1	SCSI-1	1986	~ 2 MB/sec	Asynchronous, narrow
SCSI-2	SCSI-2	1989	10 MB/sec	Synchronous, wide
SCSI-3	Split command sets, transport protocols, and physical interfaces into separate standards			
Fast-Wide	SPI/SIP	1992	20 MB/sec	
Ultra	Fast-20 annex	1995	40 MB/sec	
Ultra 2	SPI-2	1997	80 MB/sec	LVD
Ultra 3	SPI-3	1999	160 MB/sec	DT, CRC
Ultra 320	SPI-4	2001	320 MB/sec	Paced, Packetized, QAS

# SCSI standards structure



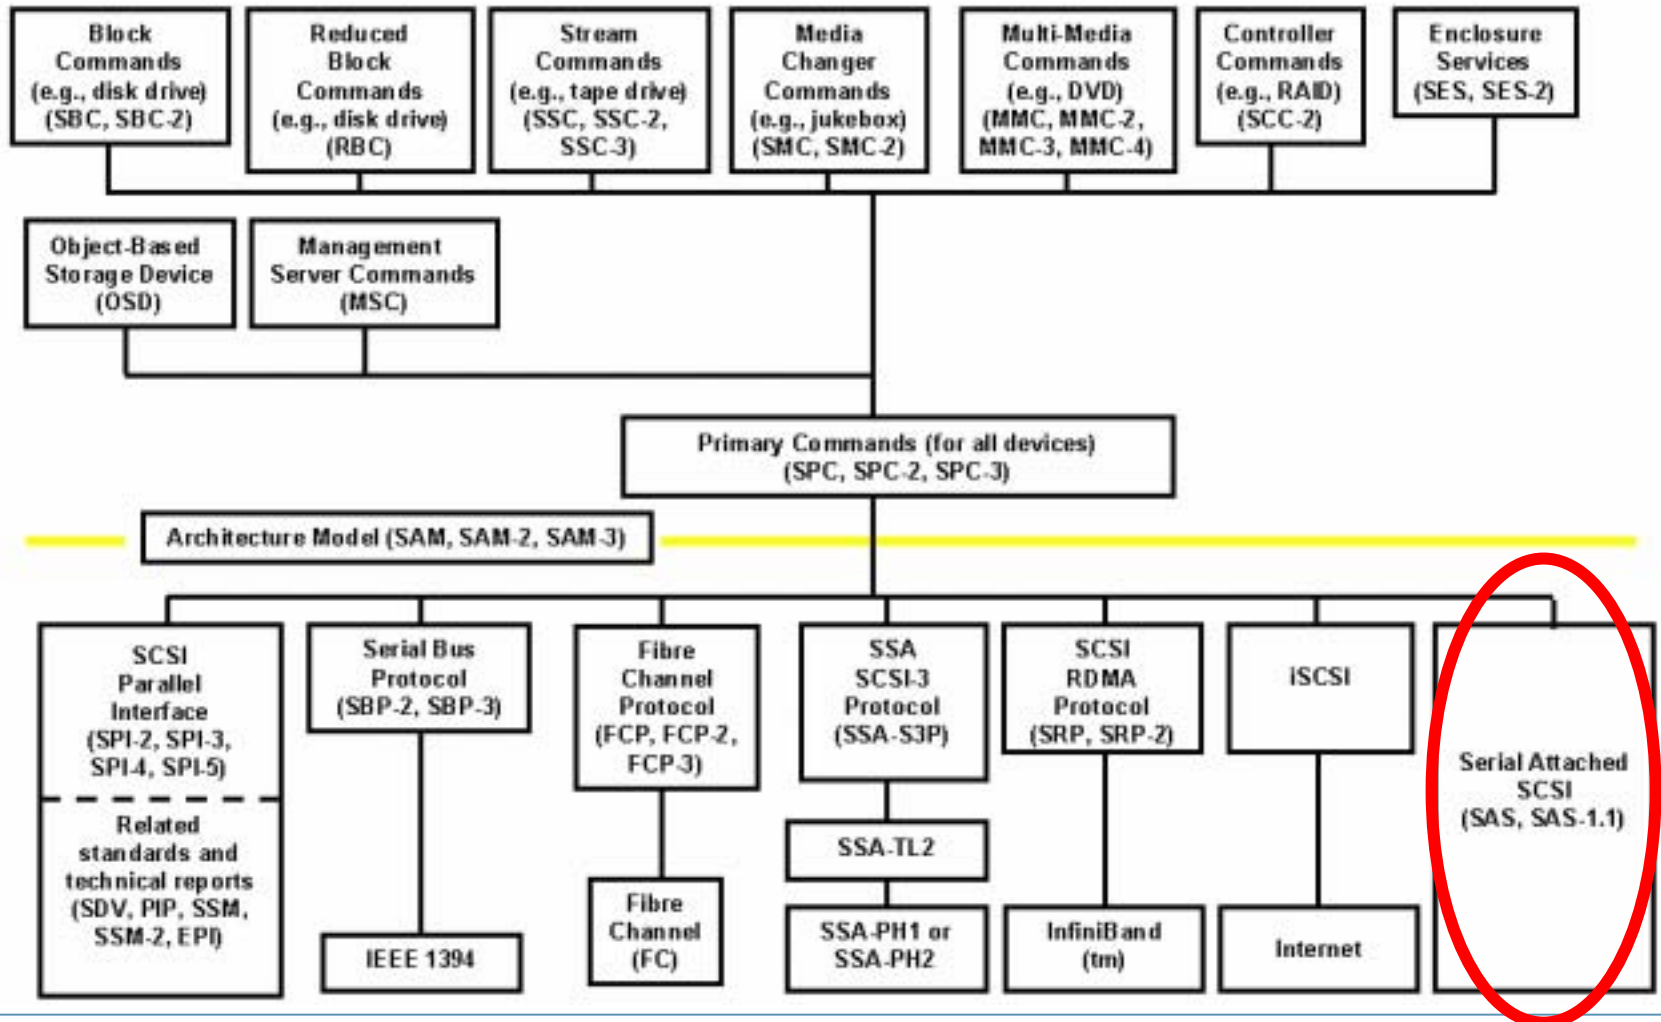
- With SCSI-3, SCSI has been broken into multiple standards (and has dropped the -3 moniker)
- Allows support for additional SCSI transport protocols



# SCSI architecture roadmap (comprehensive)



## SCSI Architecture Roadmap



# History of serial SCSI



- There have been many serial SCSI protocols

Interconnect	Standard	Year	Typical speed	Key features
Fibre Channel	FCP	1995	100 MB/sec	Optical
Serial Storage Architecture (SSA)	SSA-S2P, SSA-TL1, SSA-PL1	1996	20 MB/sec	IBM only vendor
Serial Storage Architecture (SSA)	SSA-S3P, SSA-TL2, SSA-PL2	1997	40 MB/sec	IBM only vendor
FireWire (IEEE 1394)	SBP-2	1998	50 MB/sec	
Fibre Channel	FCP-2	2002	200 MB/sec	
InfiniBand	SRP	2002	250 MB/sec	4x, 12x too
Ethernet	iSCSI	2003	~ 100 MB/sec	Gigabit Ethernet



# ATA – AT Attachment



# History of parallel ATA



Generation	Standard	Year	Speed	Key features
IDE		1986		Pre-standard
	ATA	1994		PIO modes 0-2, multiword DMA 0
EIDE	ATA-2	1996	16 MB/sec	PIO modes 3-4, multiword DMA modes 1-2, LBAs
	ATA-3	1997	16 MB/sec	SMART
	ATA/ATAPI-4	1998	33 MB/sec	Ultra DMA modes 0- 2, CRC, overlap, queuing, 80-wire
Ultra DMA 66	ATA/ATAPI-5	2000	66 MB/sec	Ultra DMA mode 3-4
Ultra DMA 100	ATA/ATAPI-6	2002	100 MB/sec	Ultra DMA mode 5, 48-bit LBA
Ultra DMA 133	ATA/ATAPI-7	2003	133 MB/sec	Ultra DMA mode 6

# ATA standards architecture



ATA register-delivered command set (ATA/ATAPI-7 Volume 1)

Device-type specific command sets (e.g., MMC-3)

Primary command set (shared for all device types) (SPC-3)

ATAPI  
(Packet delivered command set)  
(ATA/ATAPI-7 Volume 1)

ATA/ATAPI register set (ATA/ATAPI-7 Volume 1)

Protocols (e.g., SAS STP, ATA/ATAPI-7 Volume 2 (parallel ATA), ATA/ATAPI-7 Volume 3 (Serial ATA))

Interconnects (e.g., SAS, ATA/ATAPI-7 Volume 2 (parallel ATA), ATA/ATAPI-7 Volume 3 (Serial ATA))

# History of serial ATA



Generation	Standard	Year	Speed	Key features
Serial ATA	ATA/ATAPI-7	2002	150 MB/sec	
Serial ATA II	ATA/ATAPI-8?	2003?	300 MB/sec	Enhanced queuing
Serial ATA III	ATA/ATAPI-9?	?	600 MB/sec	

# Distinct markets




- Different classes of disk drives

Characteristic	Mobile	Desktop	Enterprise
<b>rpm</b>	3600, 4200, 5400 rpm	5400, 7200 rpm	10K, 15K rpm
<b>Seek time</b>	12 – 14 ms	8.9 – 9.5 ms	3.2 – 7.4 ms
<b>Performance as file server*</b>	N/A	79 – 136	146 - 366
<b>Write cache</b>	2 MB	2 – 8 MB	2 – 8 MB
<b>Capacity</b>	10 – 80 GB	40 – 250 GB	18, 36, 72, 144, 180 GB
<b>Reliability</b>	300 K hr MTBF	500 K hr MTBF	1.2 M hr MTBF
<b>Power</b>	2.5 W	10 W	15 W
<b>Cost</b>	\$73 – \$160	\$75 – \$240	\$160 – \$1400
<b>Interfaces</b>	ATA/66, ATA/100	ATA/100, ATA/133	Ultra 160 SCSI, Ultra 320 SCSI, FC

\* Benchmark of many drives on <http://www.storagereview.com>

Note: As of mid 2002



## Serial Attached SCSI

# Serial Attached SCSI (SAS)



- Serial Attached SCSI (SAS) supports both SCSI and ATA
- Three transport protocols
  - Serial SCSI Protocol (SSP)
    - Supports SAS (SCSI) disk drives, tape drives, etc.
  - Serial ATA Tunneling Protocol (STP)
    - Supports Serial ATA disk drives
  - Serial Management Protocol (SMP)
    - Supports SAS expanders


# Positioning SATA, SAS, and FC



## Desktop PC / Workstation



## Mainstream Servers / Direct Attached Storage



## Storage Area Network



Configuration Flexibility & Simplicity  
 One host design accommodates **SATA** for Low Cost Bulk Storage or **SAS** for Performance & Reliability in Mission Critical Apps

## Desktop-class disk drives (single port)

## Enterprise-class disk drives (single or dual port)

## Enterprise-class disk drives (dual port)

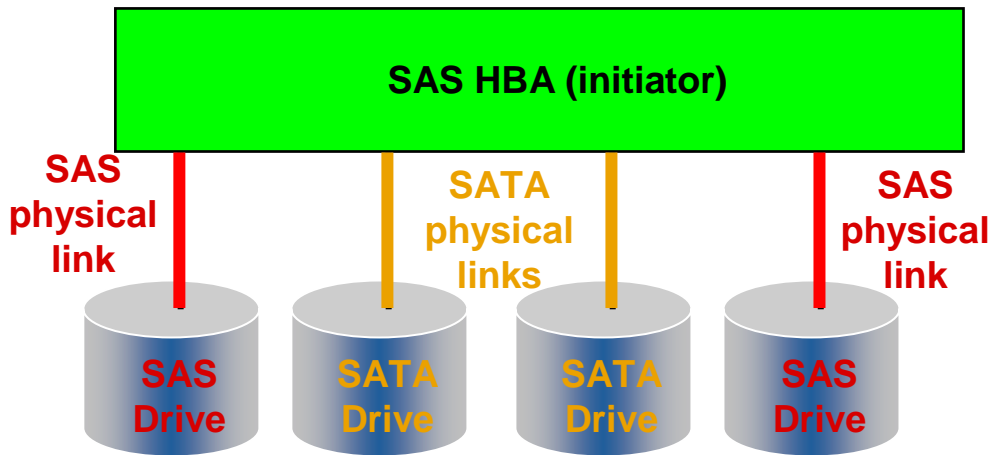
Compatible physical layer between SATA and SAS

Similar HDD ASIC and firmware design between SAS and FC

# Direct attach and expander attach



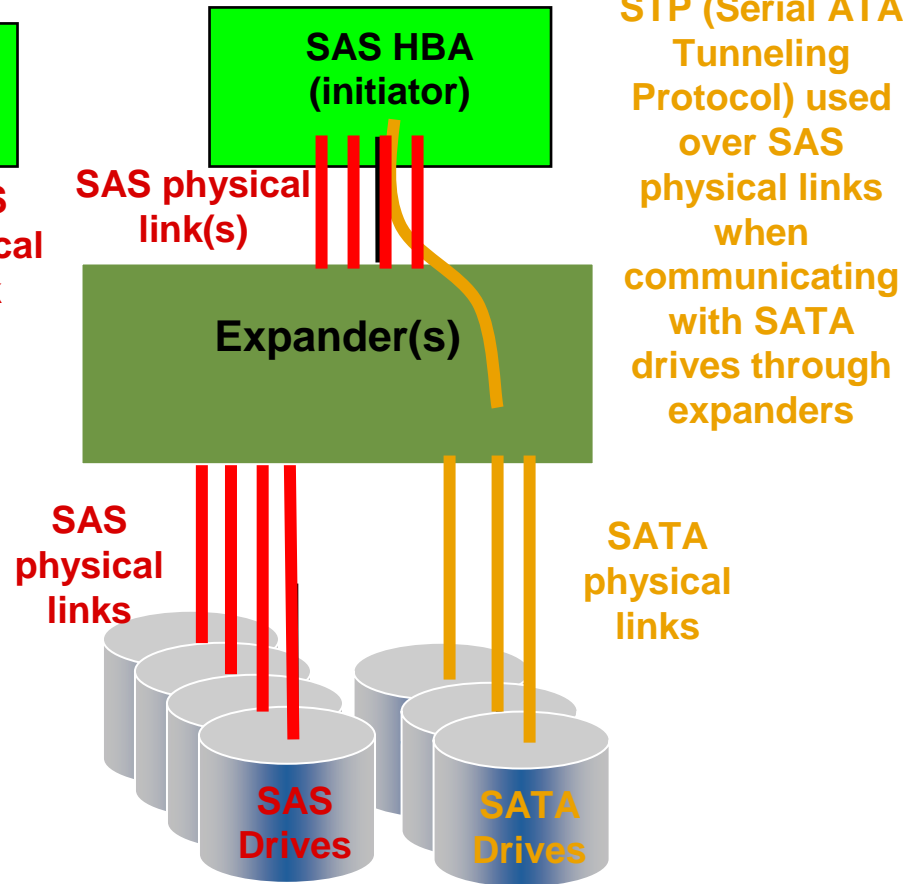
**Direct attach =**  
Number of drives limited to number of ports  
in the HBA



**SSP (Serial SCSI Protocol) used to communicate with SAS drives**

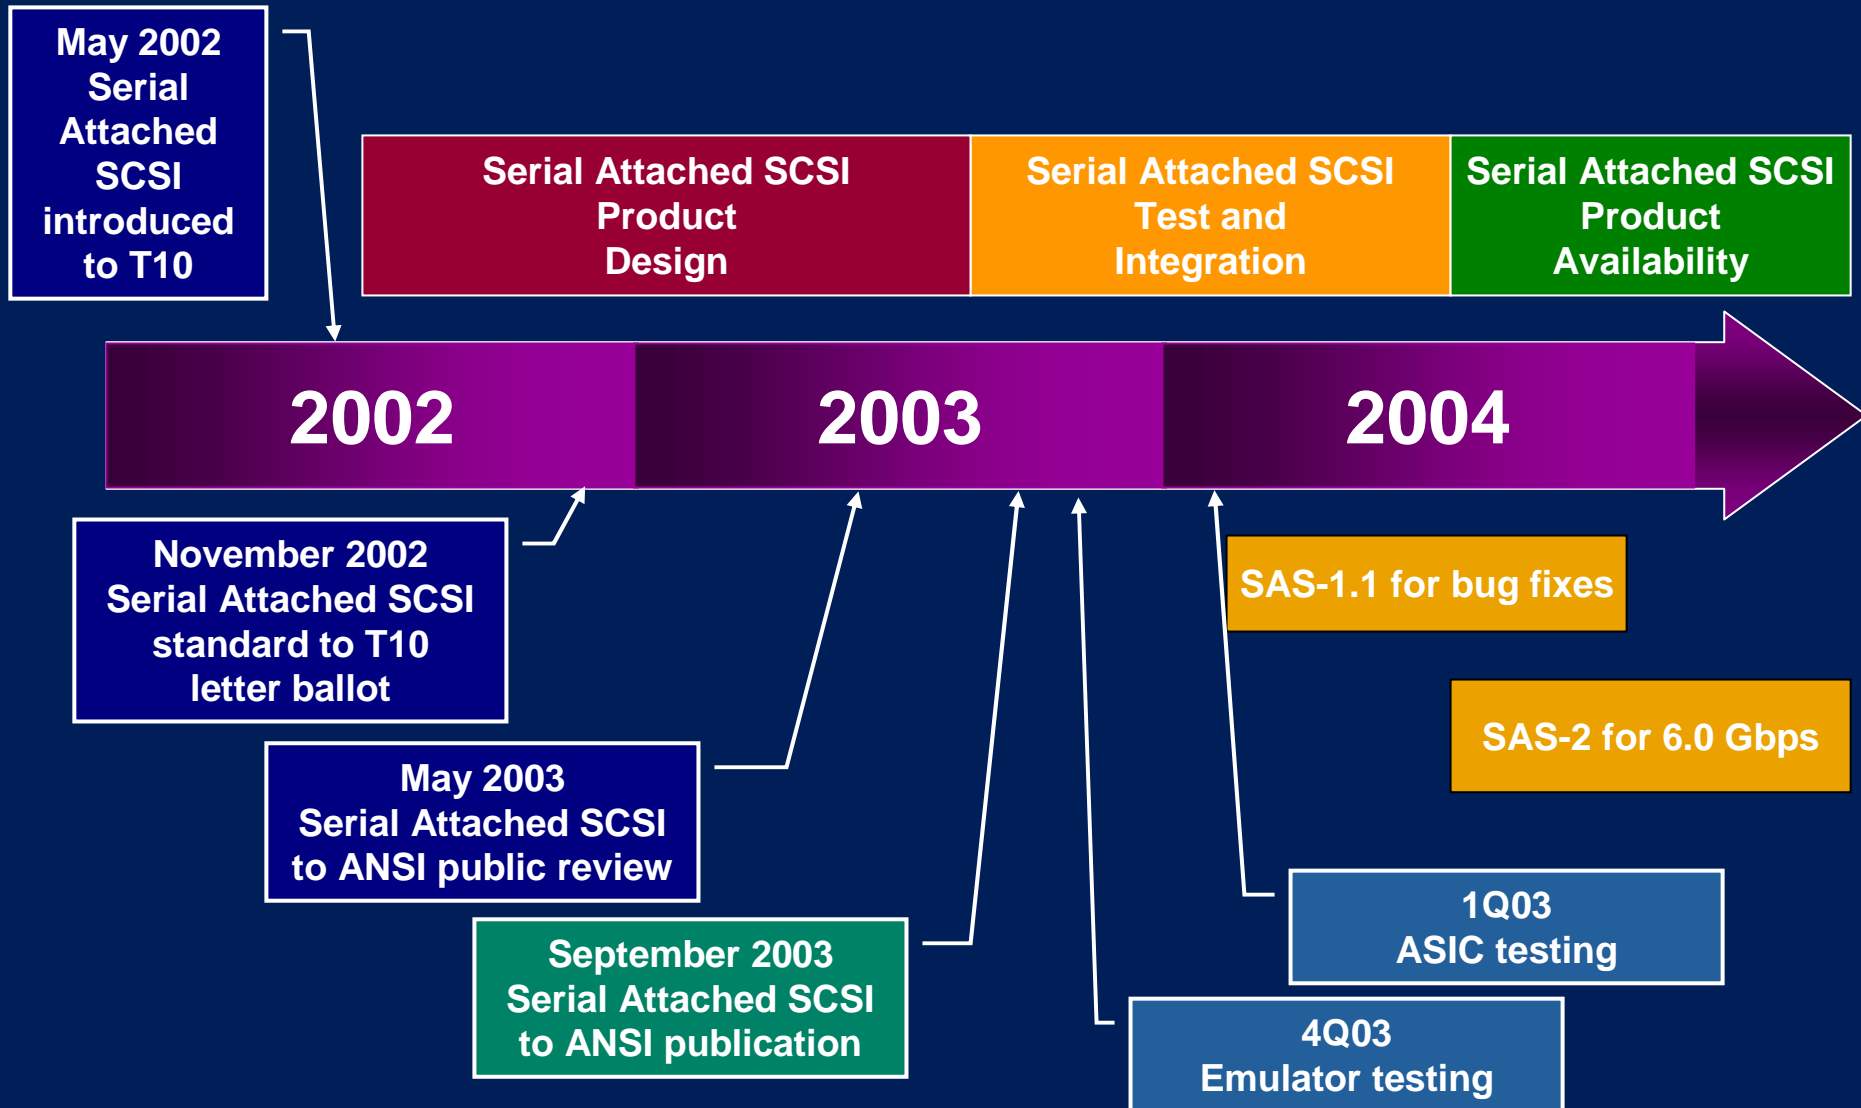
**SATA (Serial ATA) used to communicate with SATA drives over SATA physical links**

**Expander attach =**  
More drives than HBA ports






# Serial Attached SCSI timeline



# SCSI Trade Association roadmap



# SAS press releases (part 1)



- 11/26/2001 **Compaq/IBM/LSI Logic/Maxtor/Seagate**
  - Industry Leaders Team to Form Serial Attached SCSI Working Group to Address Next-Generation Interconnect Needs
- 2/4/2002 **STA**
  - SCSI Trade Association and Serial Attached SCSI Working Group Announce Agreement
- 5/6/2002 **STA**
  - INCITS Technical Committee T10 Accepts Proposal to Begin Work on Serial Attached SCSI Standard
- 6/12/2002 **Dell/HP/Intel**
  - Industry Leaders Collaborate on New Choices for Next-Generation Serial Architectures for Server Storage
- 1/9/2003 **STA**
  - Serial Attached SCSI Letter Ballot Released by T10
- 1/20/2003 **STA**
  - STA and SATA II Working Group Agree to Collaborate on Specification Alignment
- 1/20/2003 **HP/Seagate**
  - Seagate Demonstrates Industry's First Serial Attached SCSI Storage Solutions with HP
- 1/20/2003 **LSI Logic**
  - LSI Logic and Leading Hard Disk Drive Manufacturers Initiate Serial Attached SCSI Product Planning

# SAS press releases (part 2)



- 3/18/2003 **Adaptec**
  - Adaptec First Host Bus Adapter Provider to Transmit Serial Attached SCSI Protocol Packet
- 3/25/2003 **Hitachi/LSI Logic**
  - Hitachi and LSI Logic to Accelerate Delivery of Serial Attached SCSI
- 4/15/2003 **Emulex/Intel**
  - Emulex and Intel to Develop first Storage Processors for Serial ATA, SAS and Fibre Channel within a Single Architecture
- 4/21/2003 **Adaptec/Fujitsu/Hitachi/Maxtor/Seagate**
  - Adaptec Teams with Fujitsu, Hitachi, Maxtor and Seagate to Bring First Serial Attached SCSI Solutions to Market
- 4/29/2003 **LSI Logic/Seagate/Tabernus**
  - Tabernus, LSI Logic and Seagate Technology Work Together for Serial Attached SCSI
- 5/9/2003 **LSI Logic**
  - LSI Logic First To Demonstrate Serial Attached SCSI Initiator-to-Target Functionality with Full Read/Write Validation
- 6/11/2003 **Adaptec/Maxtor/Tabernus**
  - Tabernus, Adaptec, and Maxtor Work to Ensure Compatibility of Serial Attached SCSI Solutions, Speed Time to Market

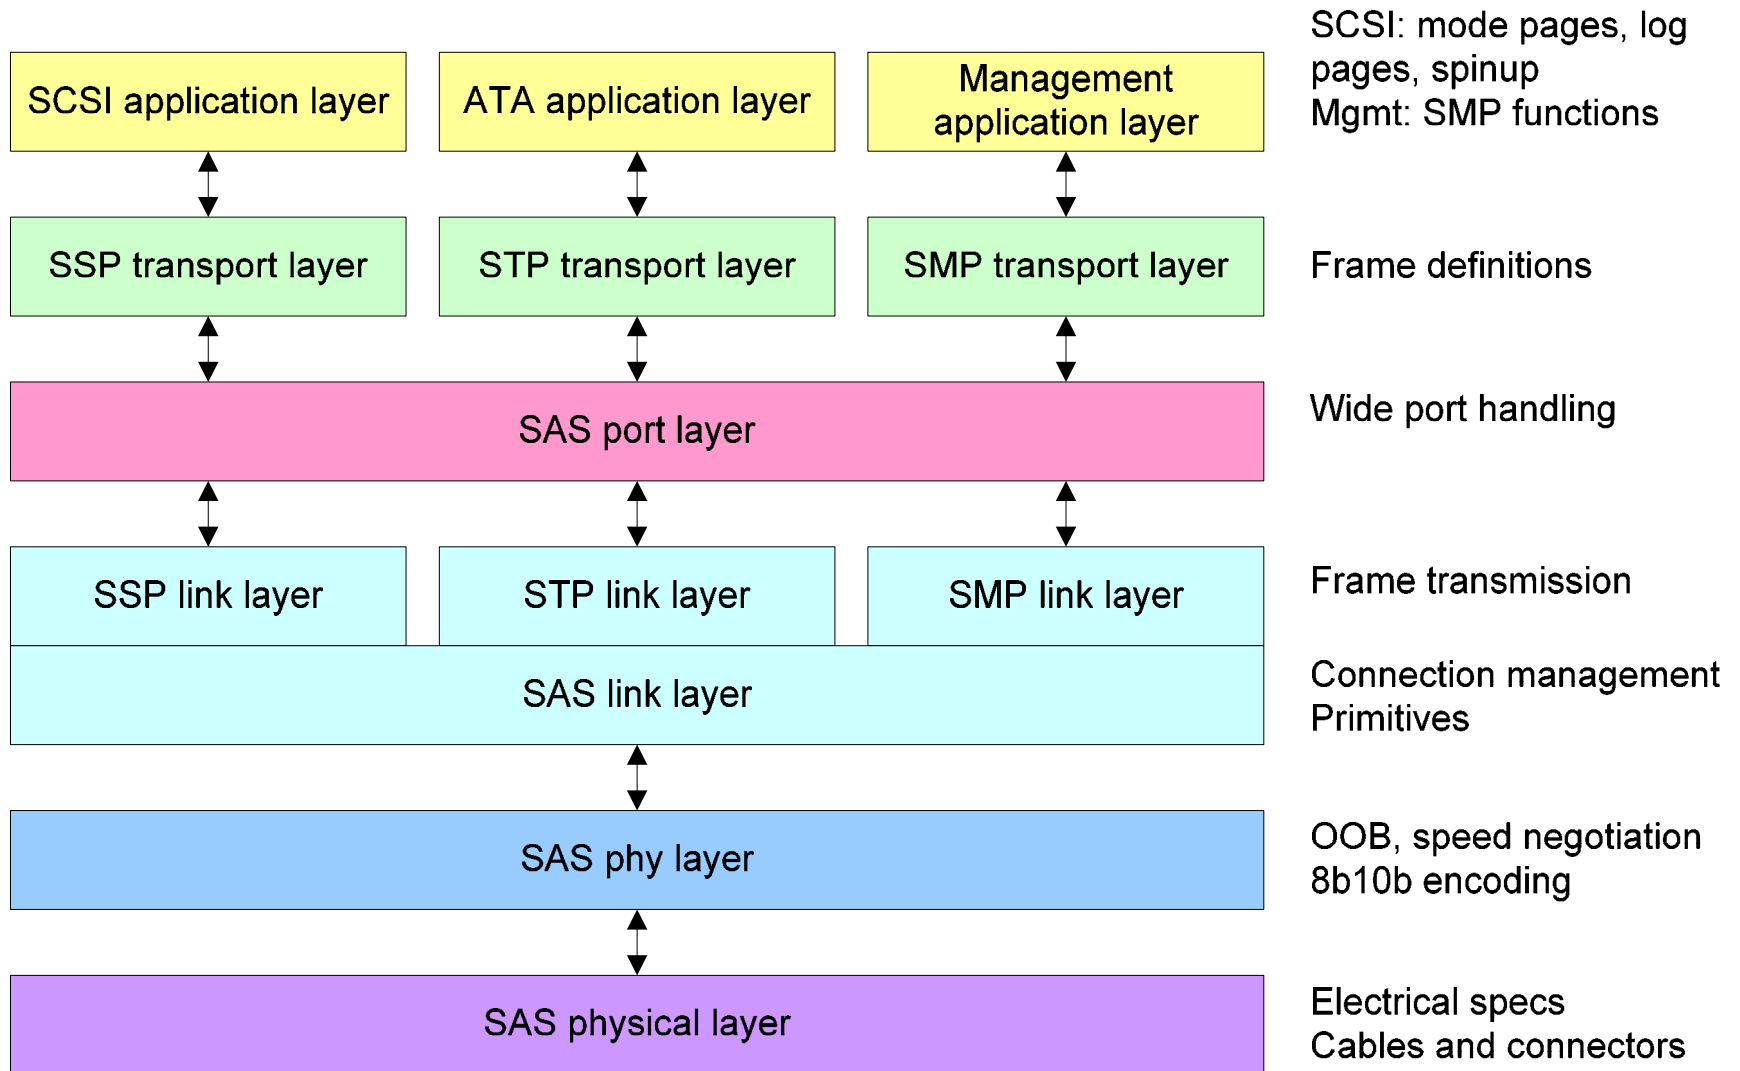
# SAS press releases (part 3)



- 6/16/2003 **Adaptec/HP/Seagate**
  - Adaptec, HP and Seagate Technology Team to Debut Compatibility of Serial Attached SCSI and Serial ATA at CeBIT America
- 6/24/2003 **LSI Logic/Maxtor**
  - LSI Logic and Maxtor Extend Testing of Serial Attached SCSI Prototypes
- 6/26/2003 **LSI/Seagate**
  - LSI Logic Teams With Seagate to Drive Serial Attached SCSI Interoperability
- 6/30/2003 **Maxtor**
  - Maxtor Demonstrates Early Success with Serial Attached SCSI Interface
- to be continued...

# Serial Attached SCSI standard

# SAS standard layering



# Key features of SAS by layer


- Physical layer – connectors
- Architecture - Expanders and topologies
- Phy layer – OOB and speed negotiation
- Link layer - connections
- Transport layer – frames
- Wrap Up




# Physical layer

# 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




**Internal backplane environment**




# SATA device connector

## Appearance of Serial ATA Connectors (Drawing courtesy of Molex)




## Device connector sizes and locations



(5.25" form factor also defined for devices like tape drives and DVDs)


in comparison...




# 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




# SAS cables and connectors - external



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


## External cabled environment



(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)




# Architecture



# Physical links and phys


- A phy contains one transceiver
- A physical link attaches two phys together



# Differential signaling




Positive signal  
(single-ended)  
(non-inverted)




Common mode  
voltage (the  
level is not very  
important)

Negative signal  
(single-ended)  
(inverted)



Common mode  
voltage


Differential signal  
positive - negative  
 $1500 - 900 = 600 \text{ mV}$   
or  
 $900 - 1500 = -600 \text{ mV}$



Differential  
signal is  
immune to  
noise common  
to both single-  
ended signals

# Physical link rate

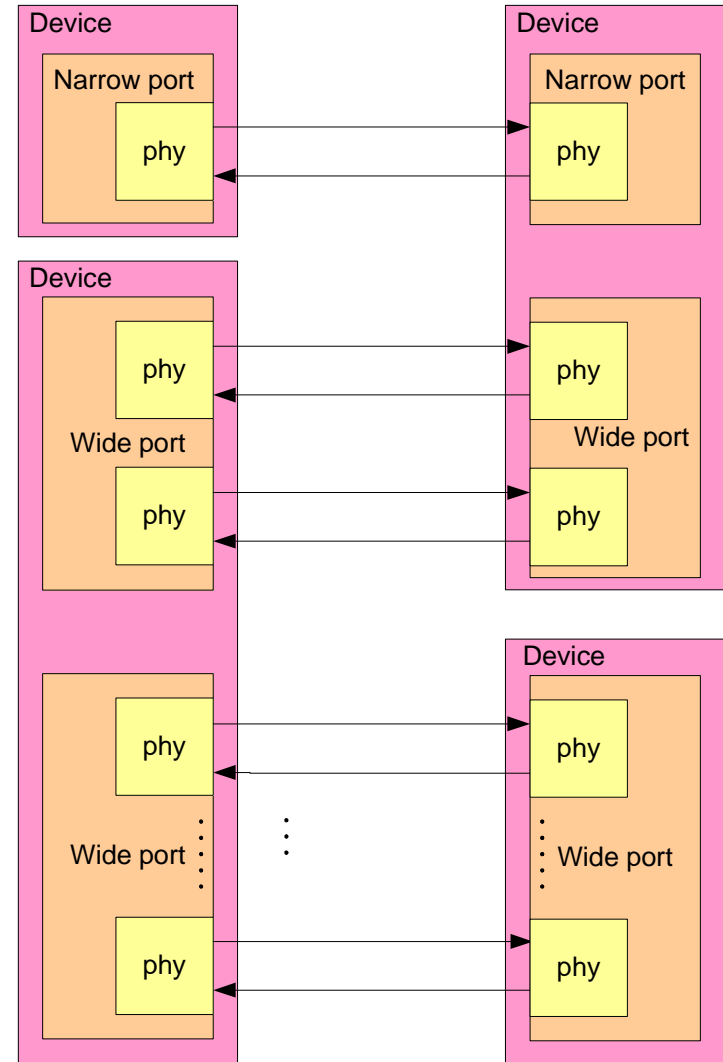
- Each direction runs 1.5 Gbps or 3.0 Gbps (150 MB/sec or 300 MB/sec)
  - Both directions use the same physical link rate
- Dual simplex (full duplex) operation – 600 MB/sec total bandwidth
- Example: peak bandwidth needs of an HBA with 8 phys
  - 2400 MB/sec half duplex, 4800 MB/sec full duplex



# SAS devices and ports



- SAS devices contain ports
- Ports contain phys
- Ports are virtual constructs
  - Groups of phy with the same SAS address, attached to another group of phys with the same SAS address



Each horizontal line represents a differential signal pair

# SAS address



- Each SAS port and expander device has a worldwide unique 64-bit SAS address
- Same namespace as the Fibre Channel Port\_Name


Byte\Bit	7	6	5	4	3	2	1	0
0	NAA (5h)							
1	IEEE Company ID (24 bits)				Vendor-Specific Identifier (40 bits)			
2								
3								
4	Vendor-Specific Identifier (40 bits)							
5								
6								
7								

- End device is a SAS device that is not an expander device
- Sample end devices
  - HBA – 8 phys
    - One SAS address for all 8 phys
      - Potentially all one (very) wide port
    - One SAS address for 4 phys, another SAS address for 4 phys
      - Guarantees at least two ports
      - Good match for 4-wide connectors
    - Eight SAS addresses
  - Disk drive - 2 phys
    - Separate SAS address for each phy
      - Guarantees two ports
      - Never a wide port


# Expander devices



- Expander device contains expander ports
- May contain SAS devices too (e.g. for enclosure management)
- Each expander device has a SAS address
- Each expander phy has a phy identifier unique within that expander device



# Expander device types – edge vs. fanout



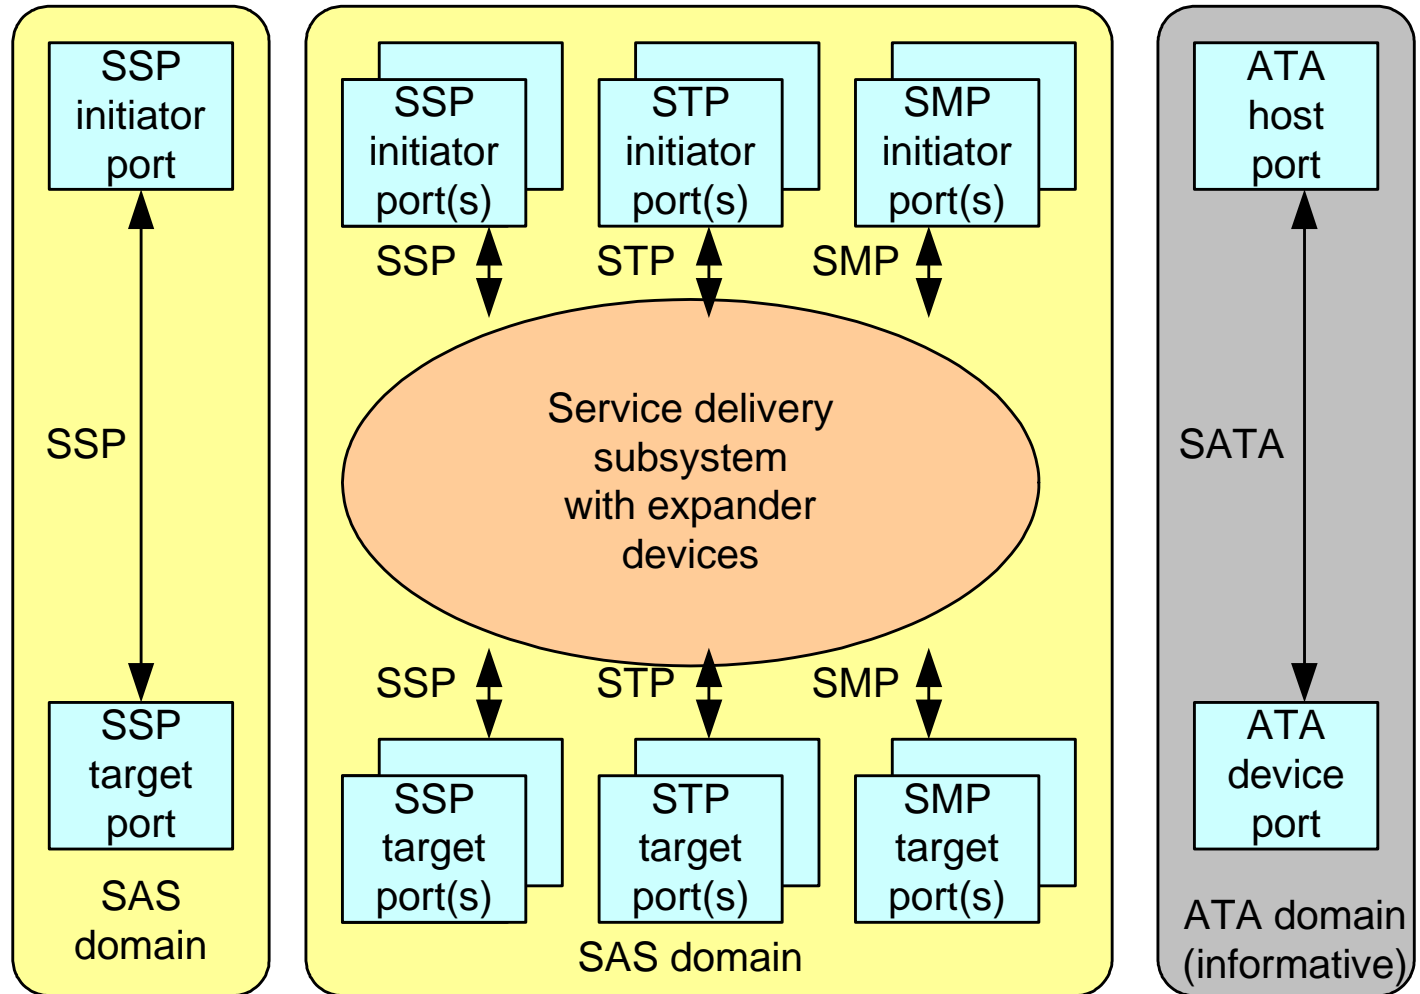
- Edge expander device
  - Always part of an “edge expander device set”
  - May perform subtractive routing
- Fanout expander device
  - Never does subtractive routing
  - Usually supports larger tables for table routing
- Topologies described later



# Domains

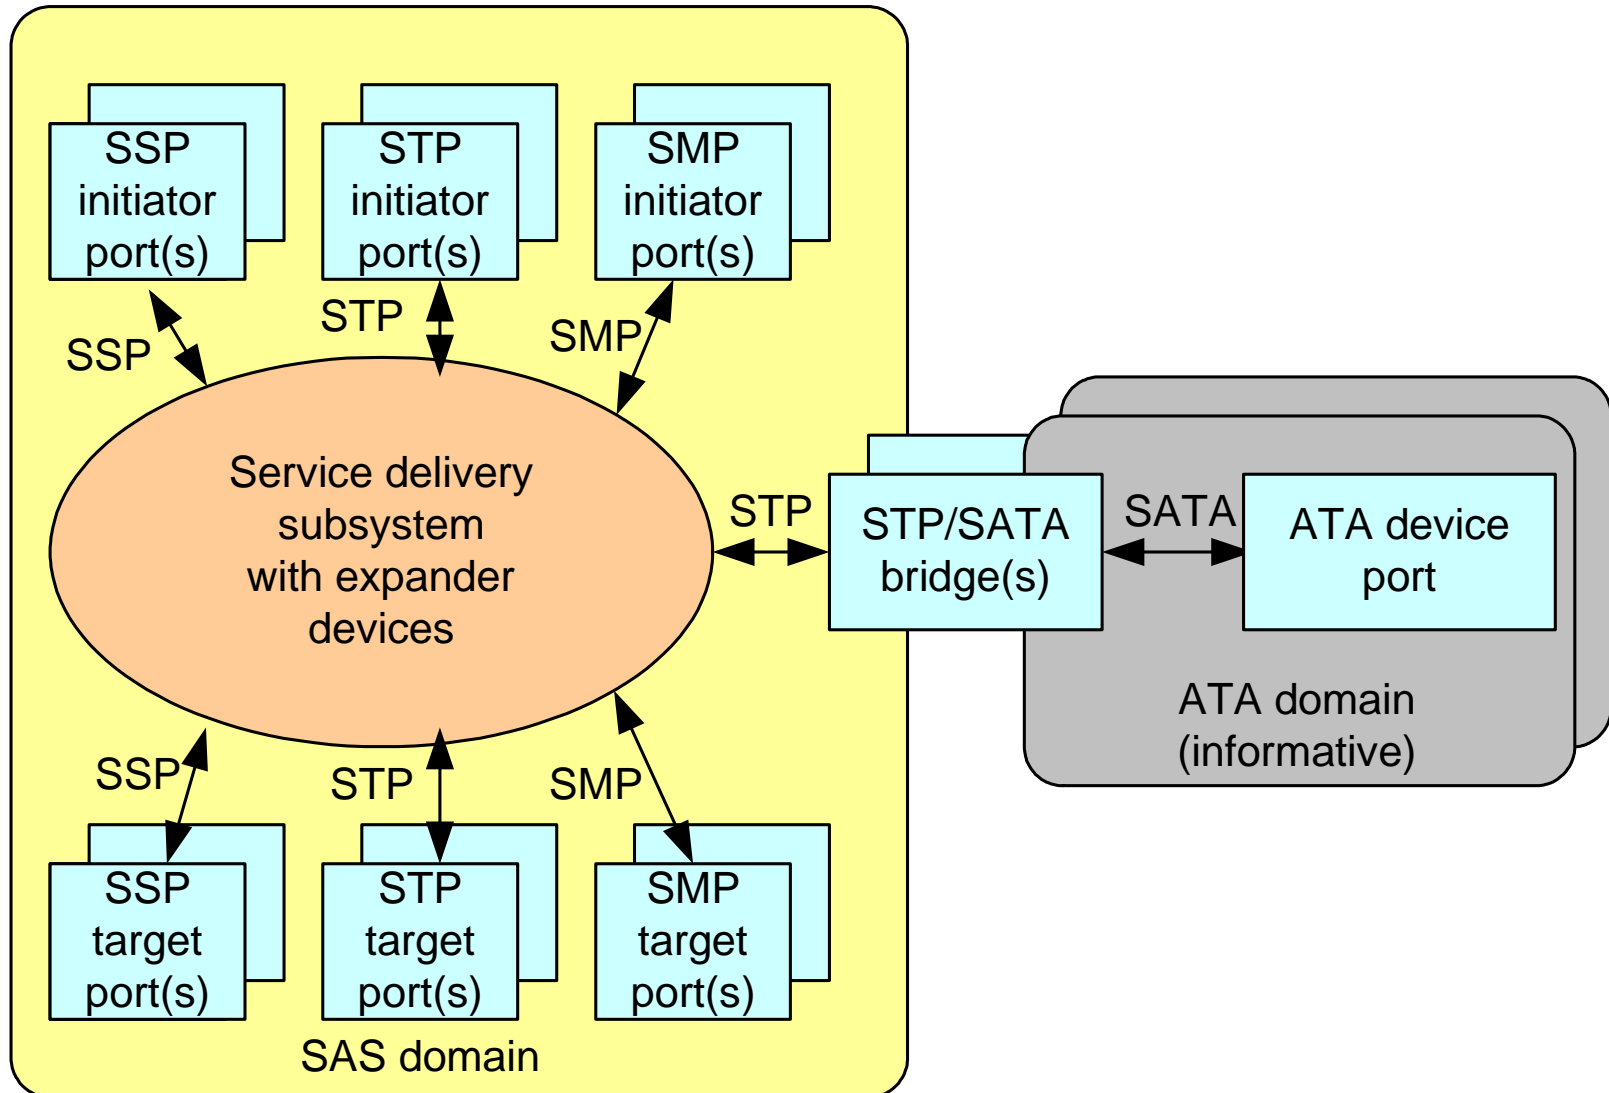


- A simple SAS domain contains SAS devices and expander devices
- An ATA domain contains a SATA host and a SATA device

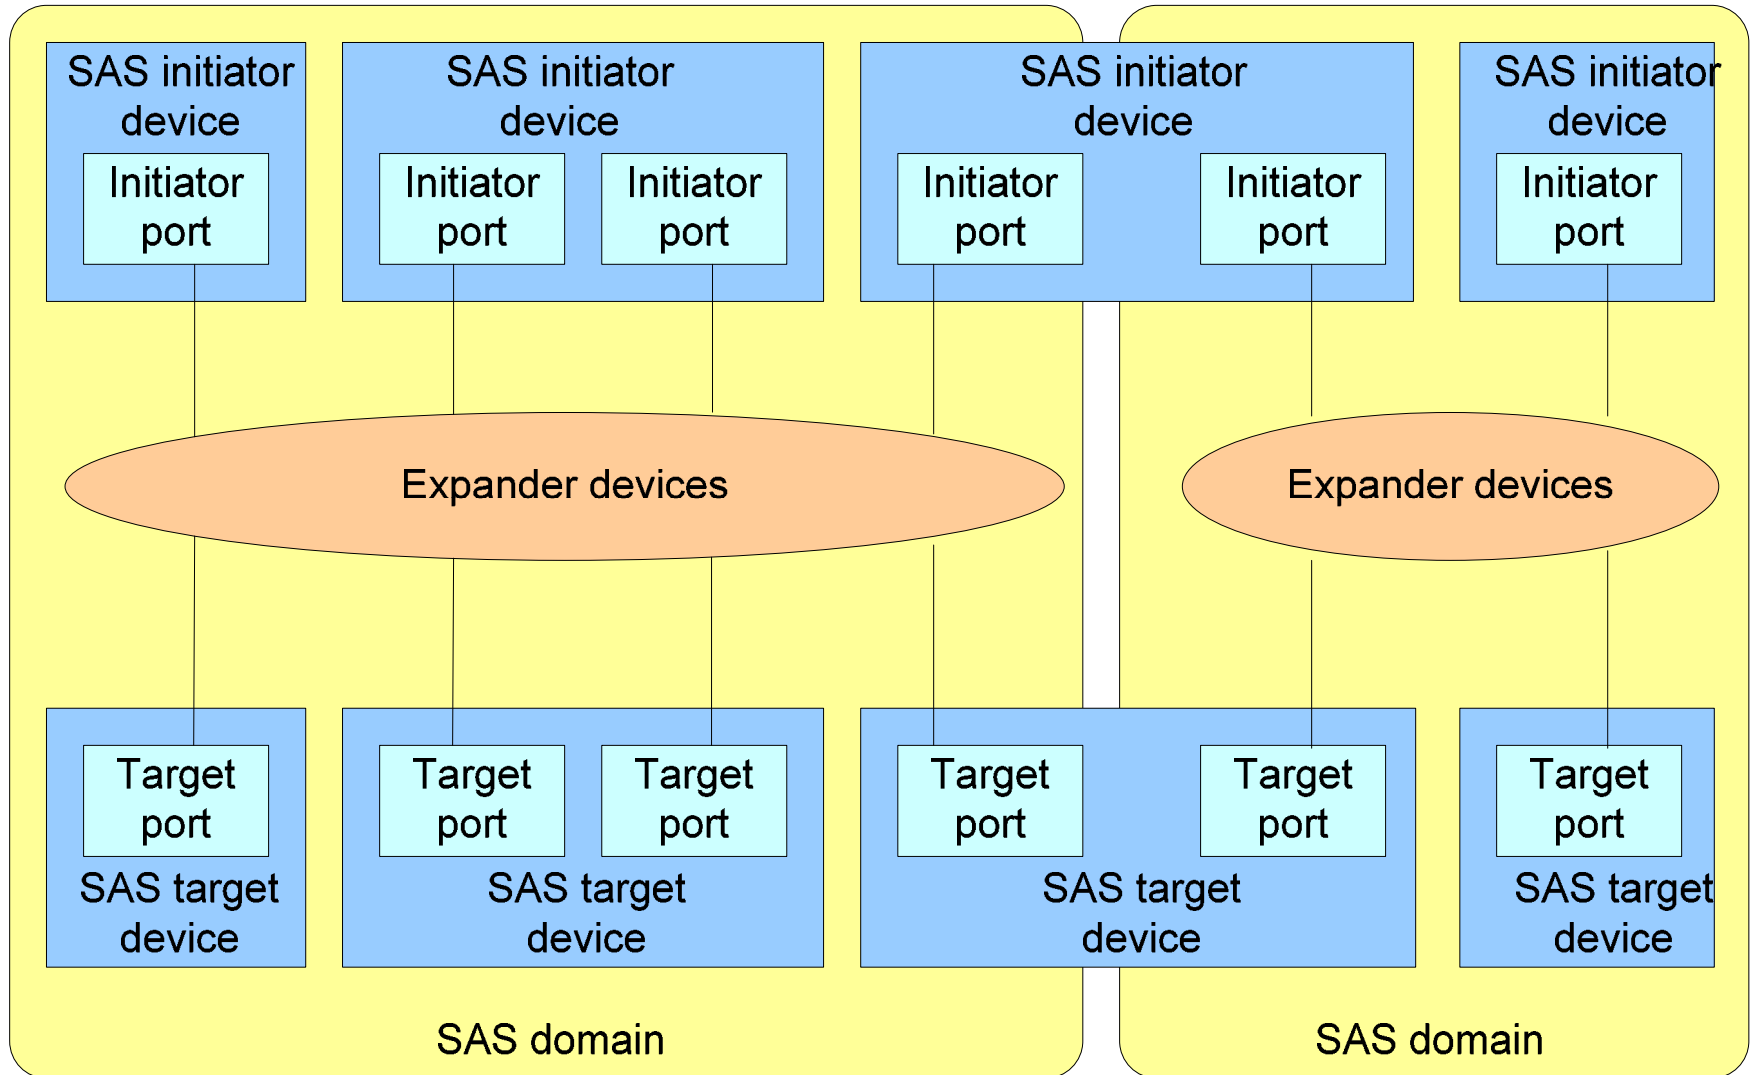


Note: When expander devices are present, SAS target ports may be located in SAS devices contained in expander devices.

# SAS domain bridged to ATA domains



# SAS devices in multiple SAS domains



# Edge expander device set



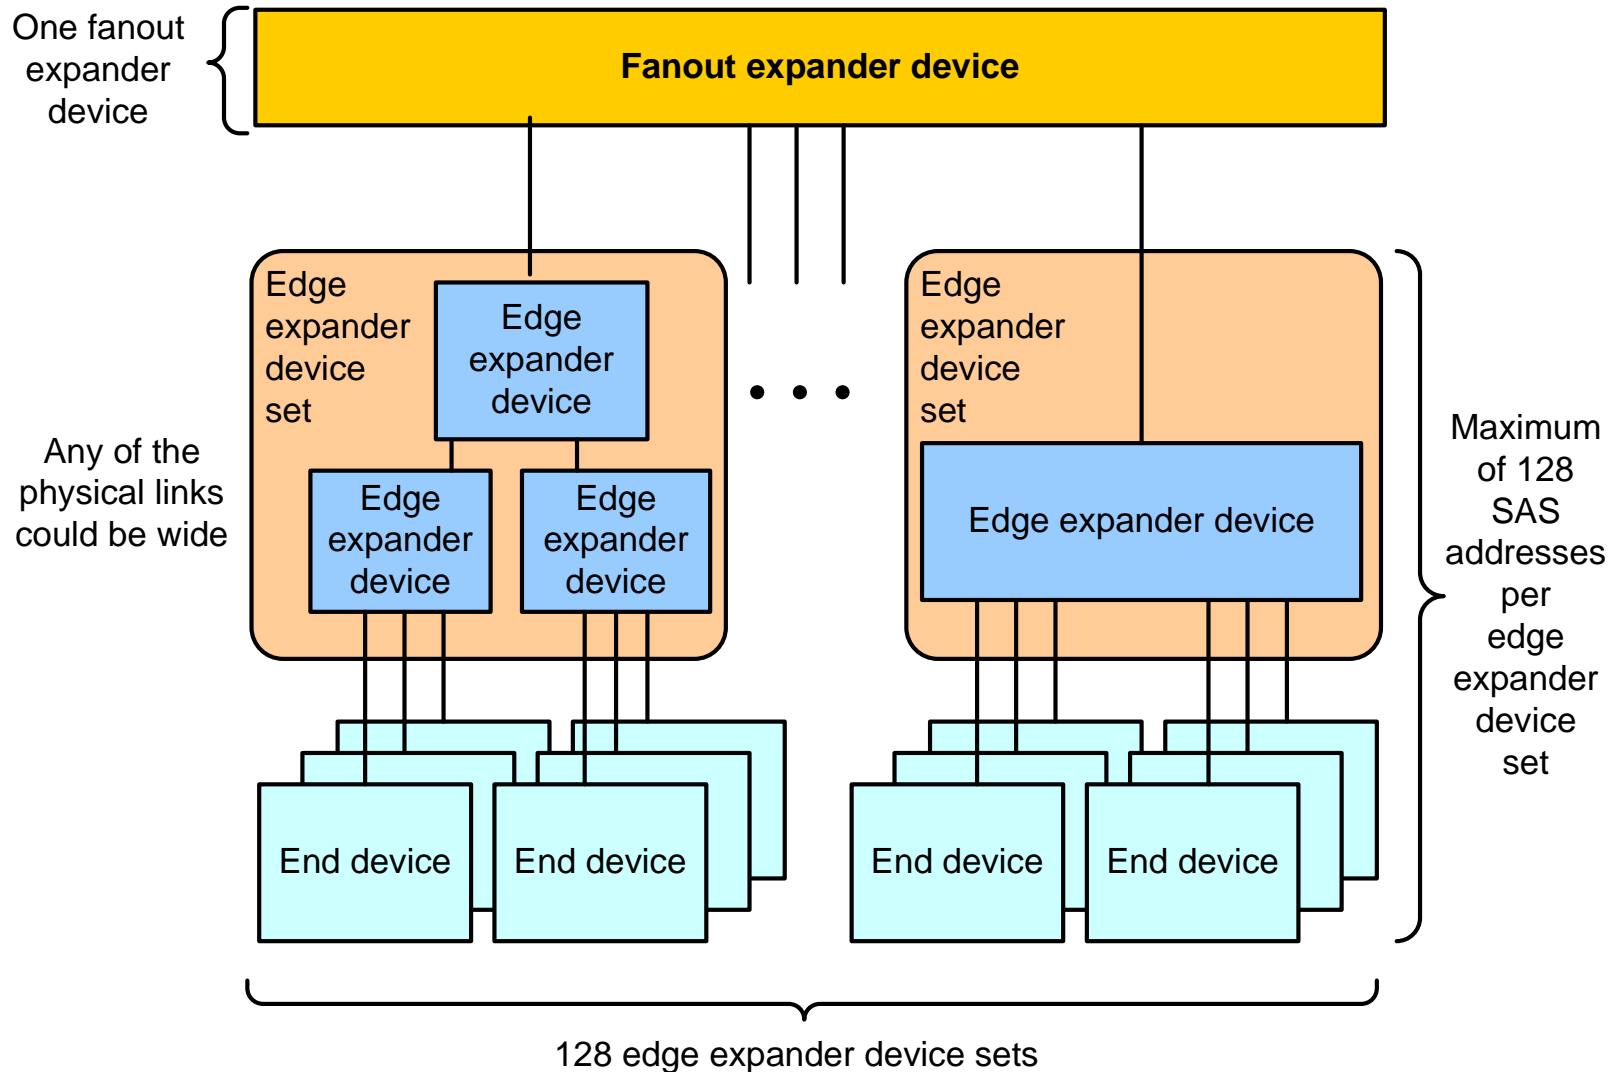
- Set of edge expander devices
- 128 SAS addresses per set
- Typically bounded by a subtractive port (to a fanout expander device, or to another edge expander device set)
- Edge expander devices uses table routing and direct routing “downstream” and subtractive routing “upstream”
- Wide links between expanders are allowed
- No loops

# Expander topologies



- Maximum of one fanout expander device in a SAS domain
- If no fanout expander, maximum of two edge expander device sets (attached via subtractive decode ports)
- End devices may be attached at any level
  - Directly to fanout expander device
  - Any level edge expander device
- Wide links possible between any two devices
- No loops
- No multiple paths

# Edge expander device set and fanout expander device




# Two edge expander device sets

The root edge expander device in each edge expander device set uses the subtractive routing method on the expander phys attached to its peer.

Any of the physical links could be wide

Upstream phys use the subtractive routing method; downstream phys use table routing method or direct routing methods.



# Expander routing attributes and methods



- Each expander phy has an expander routing attribute
- The attribute determines the routing methods the expander uses with each phy

Routing attribute	When attached to	Routing method used
Direct	End device	Direct
Table	Expander device	Table + Direct
	End device	Direct
Subtractive	Expander device	Subtractive (+ Direct)
	End device	Direct

- Direct = route requests to the attached SAS port through this phy
- Table = route requests that match in routing table through this phy
- Subtractive = route unresolved requests through this phy



# Expander route table contents



Expander route table					
	0	1	2	...	N
0					
1					
2					
...					
M					

**Phy identifier**

A phy identifier for each expander phy of the expander device.

$N \leq \text{number of phys} - 1$

**Expander route entry**

Includes:

- Routed SAS address
- Enable/disable bit


**Expander route index**

An expander route index for each expander route entry.

$M = \text{expander route indexes} - 1$

- Connection = temporary association between an initiator phy and a target phy
  - Source phy transmits an OPEN address frame
    - Contains a destination SAS address
  - Expanders route it to a matching destination phy
  - Destination phy replies with an OPEN\_ACCEPT primitive
  - Connection is established
  - Both sides exchange CLOSE primitives to close
- Connections are addressed to ports but established phy-to-phy
- N-wide ports may establish N connections at a time (to up to N other ports)
- Wide ports may establish multiple connections to other wide ports simultaneously

# Connection examples



# Connection rules

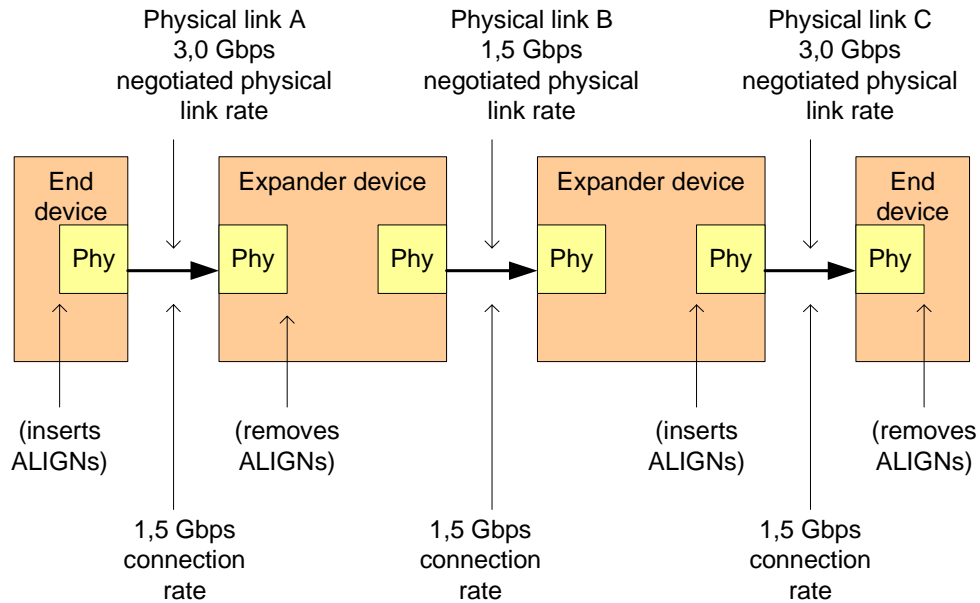


- Connections are addressed to SAS ports but are established from phy to phy
- Wide ports may establish multiple connections at a time (to up to one per phy) to different destinations
- Wide ports may establish multiple connections to other wide ports simultaneously (wide initiator port to wide target port)
  - SAS disk drives will offer two narrow ports
  - Only HBAs and RAID controllers will offer wide ports

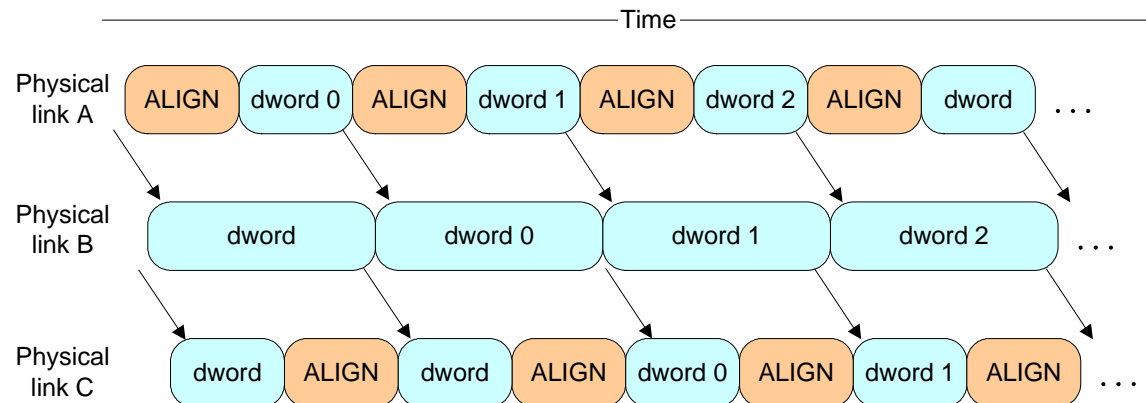
# Rate matching (connection rate)



- Lets 1.5 Gbps and 3.0 Gbps SAS ports communicate
  - Especially 3.0 Gbps initiators and 1.5 Gbps SATA drives
- On faster physical links, insert ALIGN primitives to slow down the effective data rate to match the slowest physical link in the connection



Sample dwords on physical links (from left to right) during a 1,5 Gbps connection:

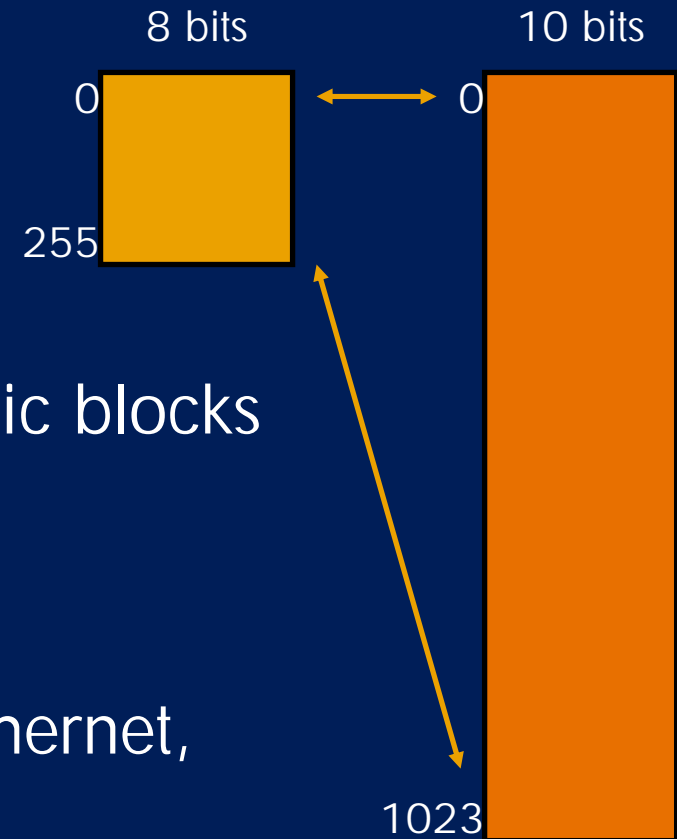


# Phy layer

# Encoding




- 8b10b coding converts 8-bit bytes into 10-bit data characters for transmission on the wire
- Reasons
  - Clock recovery
  - DC balance
  - Special characters
  - Error detection
- Mapping done with two simple logic blocks
  - 5b6b and 3b4b
  - Full table in SAS and SATA
- Invented by IBM in 1983
- Used by Fibre Channel, Gigabit Ethernet, 1394b, and many other standards



# Dwords



- **Byte** = 8 bits (xxh)
- **Character** = 10 bits (Dxx.y or Kxx.y)
- **Dword** = 4 characters (or 4 bytes, depending on context)
  - 40 bits on the wire
  - Usually represents 32 bits of data
  - A dword may flip disparity or leave it the same
  - A dword is either a **data dword**, a **primitive**, or an **invalid dword**





- 1 control character and 3 data characters
  - First character is K28.5 (for SAS primitives), K28.3 (for SATA primitives), or K28.6 (special SATA error primitive)
    - K28.6 primitive serves as an invalid dword for SATA
  - Last three characters are data characters
  - Endianness does not matter
    - both SAS and SATA primitives always have the control character first on the wire

first	second	third	fourth
K28.5	Dxx.y	Dxx.y	Dxx.y
K28.3	Dxx.y	Dxx.y	Dxx.y
K28.6	Dxx.y	Dxx.y	Dxx.y

# OOB signals

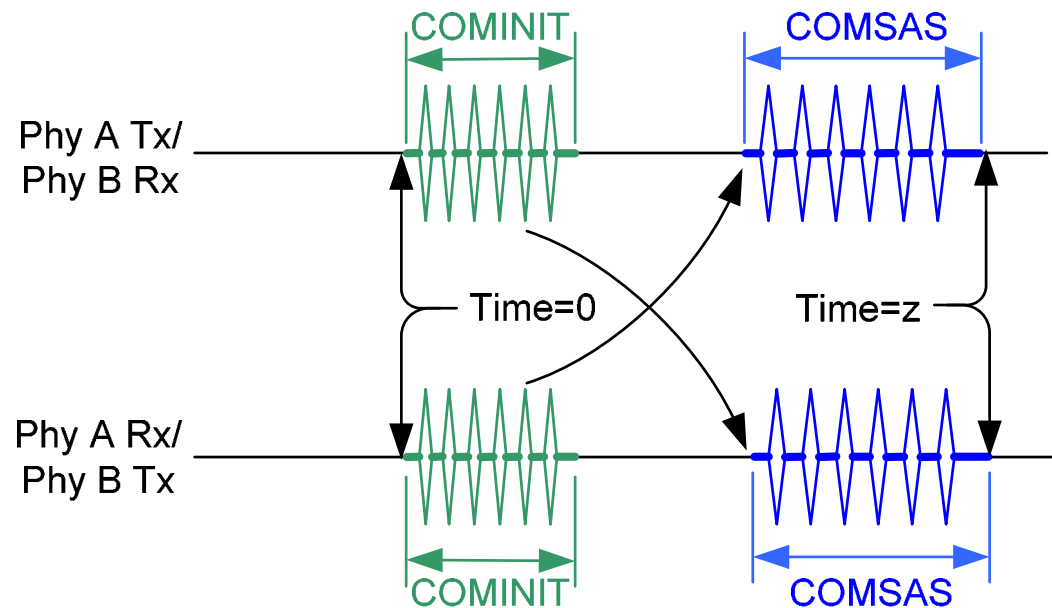


- OOB signal is a pattern of idle times and burst times
  - Idle time (and negation time)
    - Differential 0 V (Positive signal = negative signal)
    - No transitions (DC idle)
  - Burst time
    - Transmitted as a burst of ALIGN(0) primitives
    - Received as presence of edges (whether they are valid ALIGNs is irrelevant)
- Designed to be detectable by analog squelch detection logic
- Length of idle time distinguishes between OOB signals
  - COMINIT, COMWAKE, and COMSAS

# SAS OOB sequence



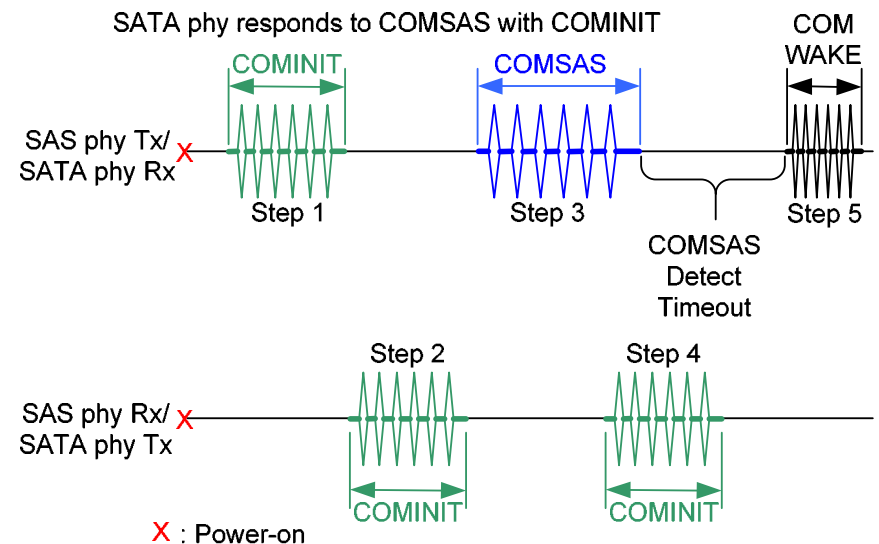
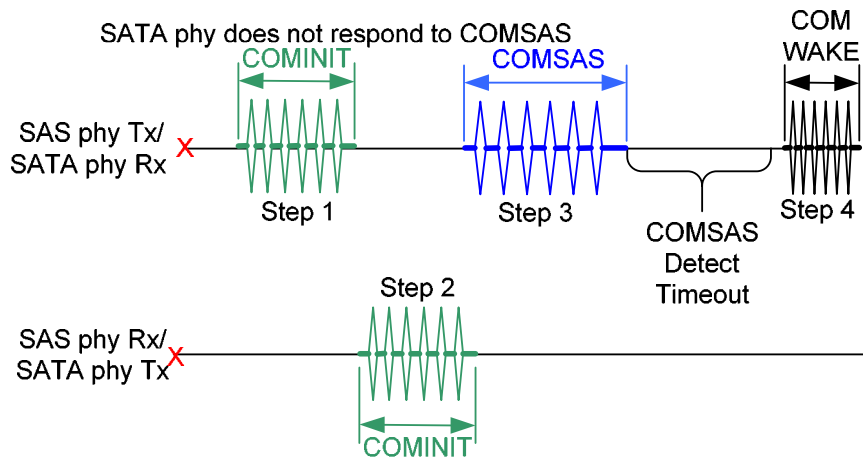
- Transmit and receive COMINIT
- Transmit COMSAS
  - If COMSAS received, physical link is SAS to SAS
  - If COMSAS is not received, physical link is SAS to SATA
- There are a few hot-plug situations where COMINIT leads directly to COMSAS; this is allowed



# SAS to SATA OOB sequence




- If COMSAS is not received, physical link is SAS to SATA
- SAS OOB sequence morphs into the SATA OOB sequence
- Since the SAS device drives COMWAKE, it looks like a SATA host
- Only supports SATA device not SATA host



# SAS speed negotiation sequence



- Slow-to-fast
- Both phys run same set of speed negotiation windows
- 1.5, then 3.0, then 6.0 (if needed), etc. until they find:
  - a supported rate; then
  - a (faster) non-supported rate
- Last window returns to the highest supported rate detected



# Link layer

# Clock skew management




- Slow down the transmitter by including dwords that can be deleted
  - ALIGN (and NOTIFY for SAS) primitives stuck in
  - Throttles the transmitter to be slower than the slowest receiver
    - SATA (and SAS STP) requires 2 per 256 dwords (0.7812%)
      - Must always appear in pairs
    - SAS requires 1 per 2048 dwords (0.0488%)
  - Receiver throws away incoming ALIGNs (and NOTIFYs for SAS)



- Established between a SAS initiator phy and a SAS target phy
- Three types of connections
  - SSP (Serial SCSI protocol)
  - SMP (Serial management protocol)
  - STP (Serial ATA Tunneling protocol)
- Basic sequence
  - 1. Transmit OPEN address frame
  - 2. Receive AIP (arbitration in progress) primitives if expanders are involved
  - 3. Receive final result
    - OPEN\_ACCEPT primitive – connection established
    - OPEN\_REJECT primitive – rejected, go back to idle
- SATA has no concept of connections



# Connection – basic sequence



# Arbitration fairness and deadlocks



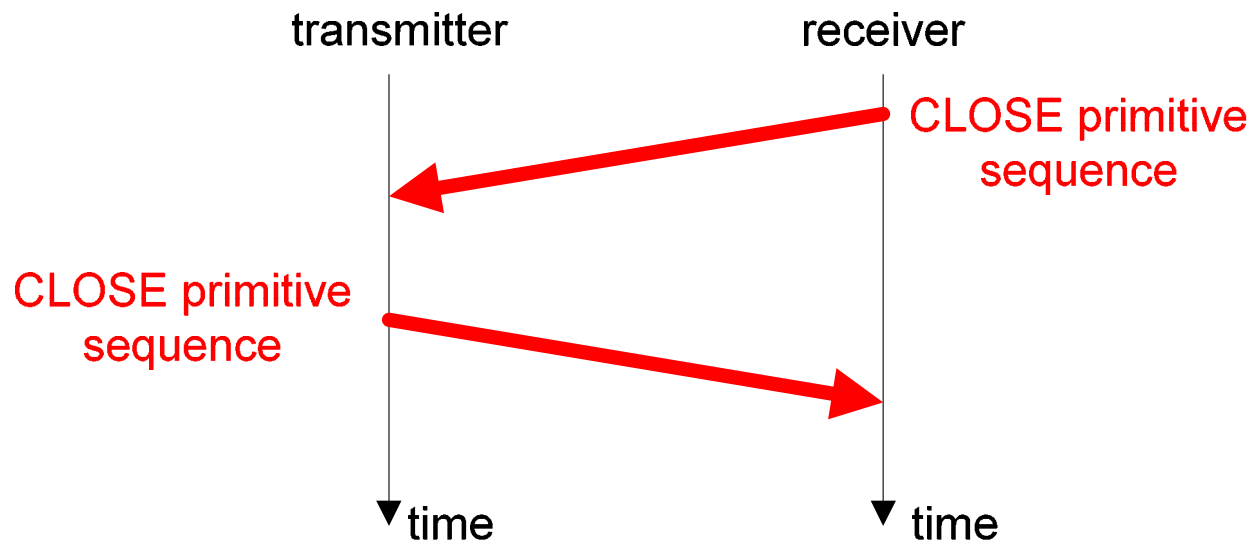
- SAS implements a Least Recently Used (LRU) arbitration fairness scheme
  - Each OPEN address frame includes the age of a request
  - In case of ties, SAS addresses are used to pick the highest priority
- SAS implements deadlock detection and recovery using Partial Pathway timers
  - When expanders detect a hang, one request is backed off to break the likely deadlock
  - SAS addresses are eventually used to decide which request is the lowest priority and gets backed off

# Closing a connection




- Exchange CLOSE primitive sequences end-to-end to close a connection
- SSP and SMP require traffic be stopped first
  - CLOSE ignored by SAS phy state machines if not at the correct time
- STP connection should be idle before closing
  - Not enforced by SL state machine
- Expanders always honor CLOSE regardless of connection state

Example 1: CLOSEs sent one at a time



Example 2: CLOSEs sent simultaneously

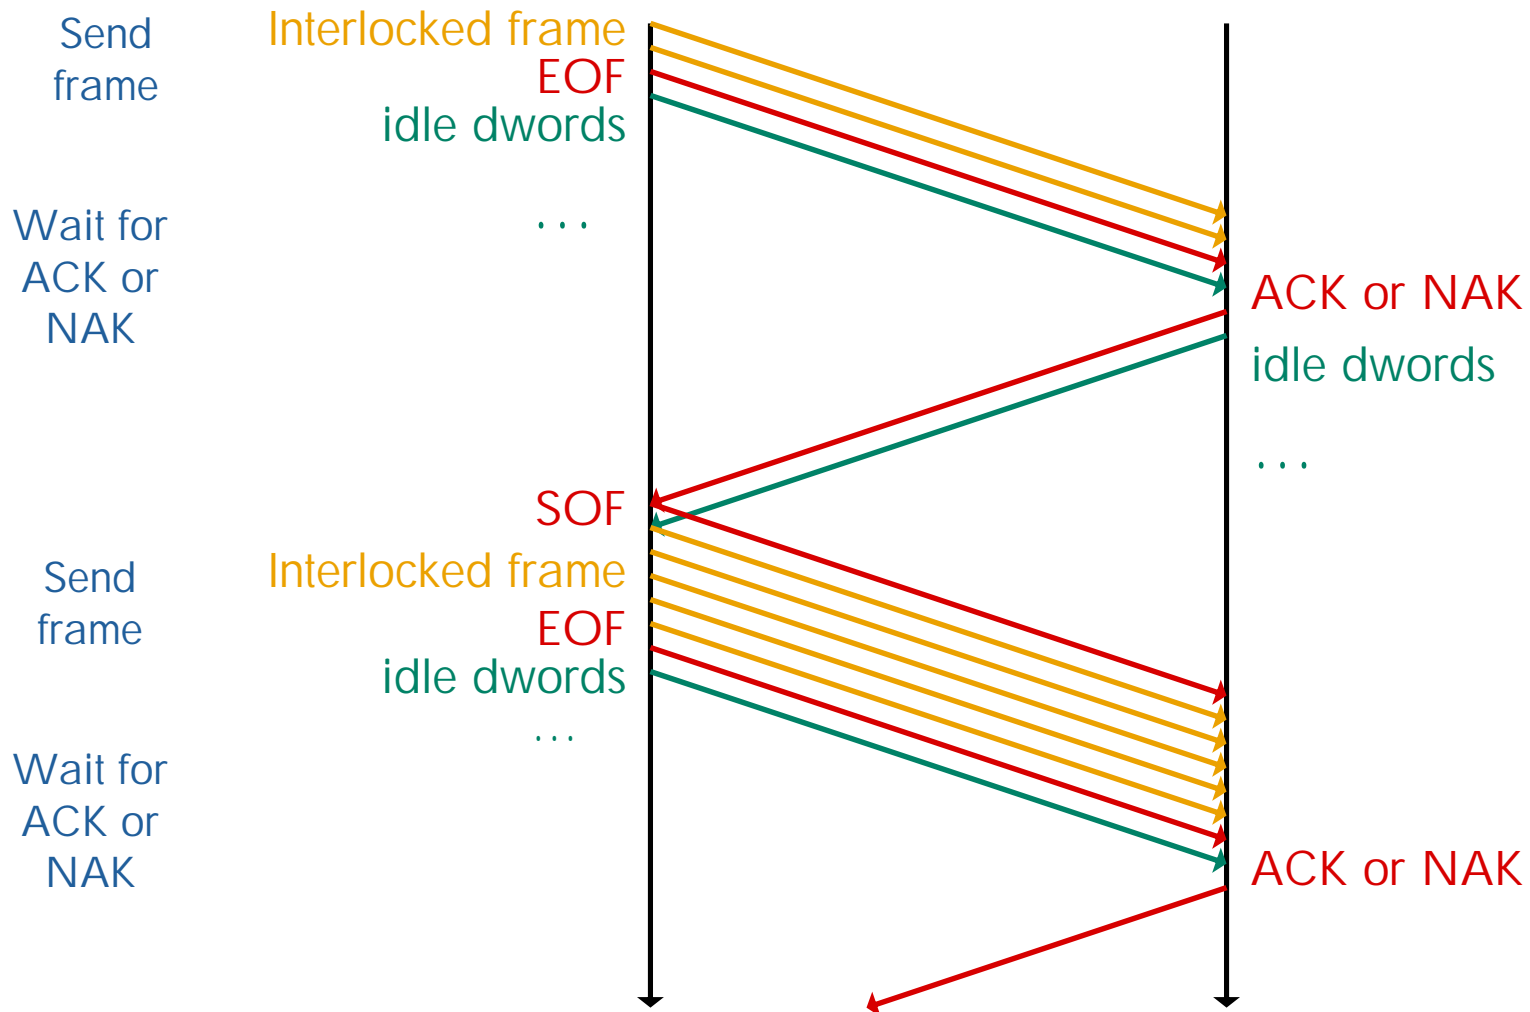


- Inside an SSP connection...
- Phys exchange **SSP frames**
  - SSP frame = SOF primitive, data dwords, EOF primitive
  - Each frame results in an ACK or NAK primitive
- Credit-based flow control
  - Permission to send a frame must be granted with RRDY primitives
- Full duplex
  - SSP frames can be sent in both directions simultaneously
  - Independent credit for each direction
- ACK, NAK, and RRDY primitives may be interjected among frame dwords
  - (so can ALIGNs and NOTIFYs)

# SSP interlocked frames example




- Showing frames only in one direction



# SSP non-interlocked frames example



- Showing frames in one direction




# SMP overview



- Only an initiator can open an SMP connection
  - Target not allowed to open an initiator
- Inside an SMP connection...
- Two SMP frames are transferred
  - SMP frame = SOF primitive, data dwords, EOF primitive
  - 1. Initiator transmits one SMP\_REQUEST frame to target
  - 2. Target transmits one SMP\_RESPONSE frame to initiator

# SMP example






# STP and SATA overview



- In SATA, SATA host and SATA device just communicate directly – no connections
- In SAS, once an STP connection is open, STP initiator and STP target communicate as if they were SATA host and SATA device directly attached on a physical link
  - Extra latency introduced as dwords flow through expanders
- Half duplex
  - SATA never transmit frames in both directions at one time
  - Usually the frame goes one way and R\_IP primitives go the other way
- Note: SAS prefixes the SATA primitive names with “SATA\_”
  - SATA\_SOF (used in STP) is different than SOF (used in SSP)
  - In this presentation, prefix not always used

# SATA basic frame transmission



# Transport layer

# SSP frames



- One common SSP frame format
  - **Frame header**: 24 bytes
  - **Information Unit**: 0 to 1024 bytes
  - **Fill bytes**: 0 to 2 bytes
  - **CRC**: 4 bytes
- Frame format is based on Fibre Channel

Byte	Field(s)		
0	Frame Type		
1 to 3	Hashed Destination SAS address		
4	Reserved		
5 to 7	Hashed Source SAS address		
8 to 9	Reserved		
10	Reserved	Retransmit	Rsvd
11	Reserved	Number of Fill Bytes	
12 to 15	Reserved		
16 to 17	Tag		
18 to 19	Target Port Transfer Tag		
20 to 23	Data Offset		
24 to m	Information Unit		
m to (n-3)	Fill bytes, if needed		
(n-3) to n	CRC		


# SSP frame types




- Information Units are nearly identical to Fibre Channel Protocol (FCP)
  - Separate TASK frame is new to SAS

Command	Information Unit field size	Direction	Description
COMMAND	28 to 284	I to T	Send a command
TASK	28	I to T	Send a task management function
XFER_RDY	12	T to I	Request write data
DATA	1 to 1024	I to T or T to I	Write data (I to T) or read data (T to I)
RESPONSE	24 to 1024	T to I	Send SCSI status (for commands) or task management response (for task management functions)

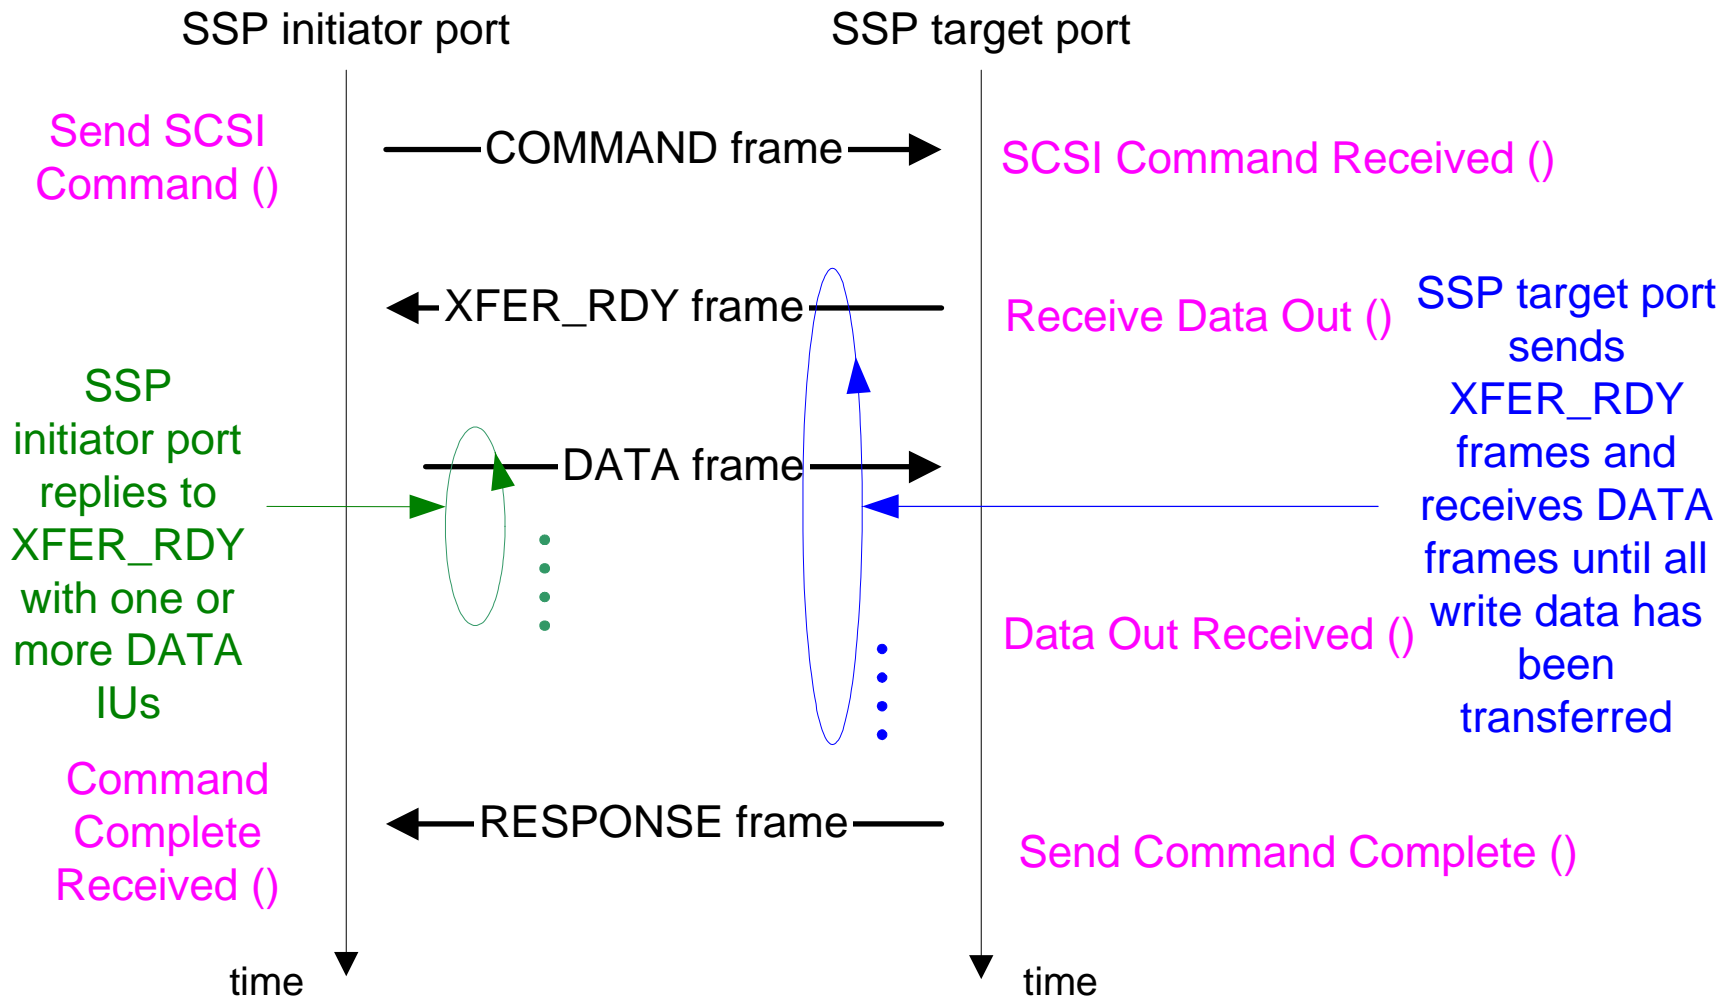
# SSP task management frame sequence



# SSP non-data command sequence

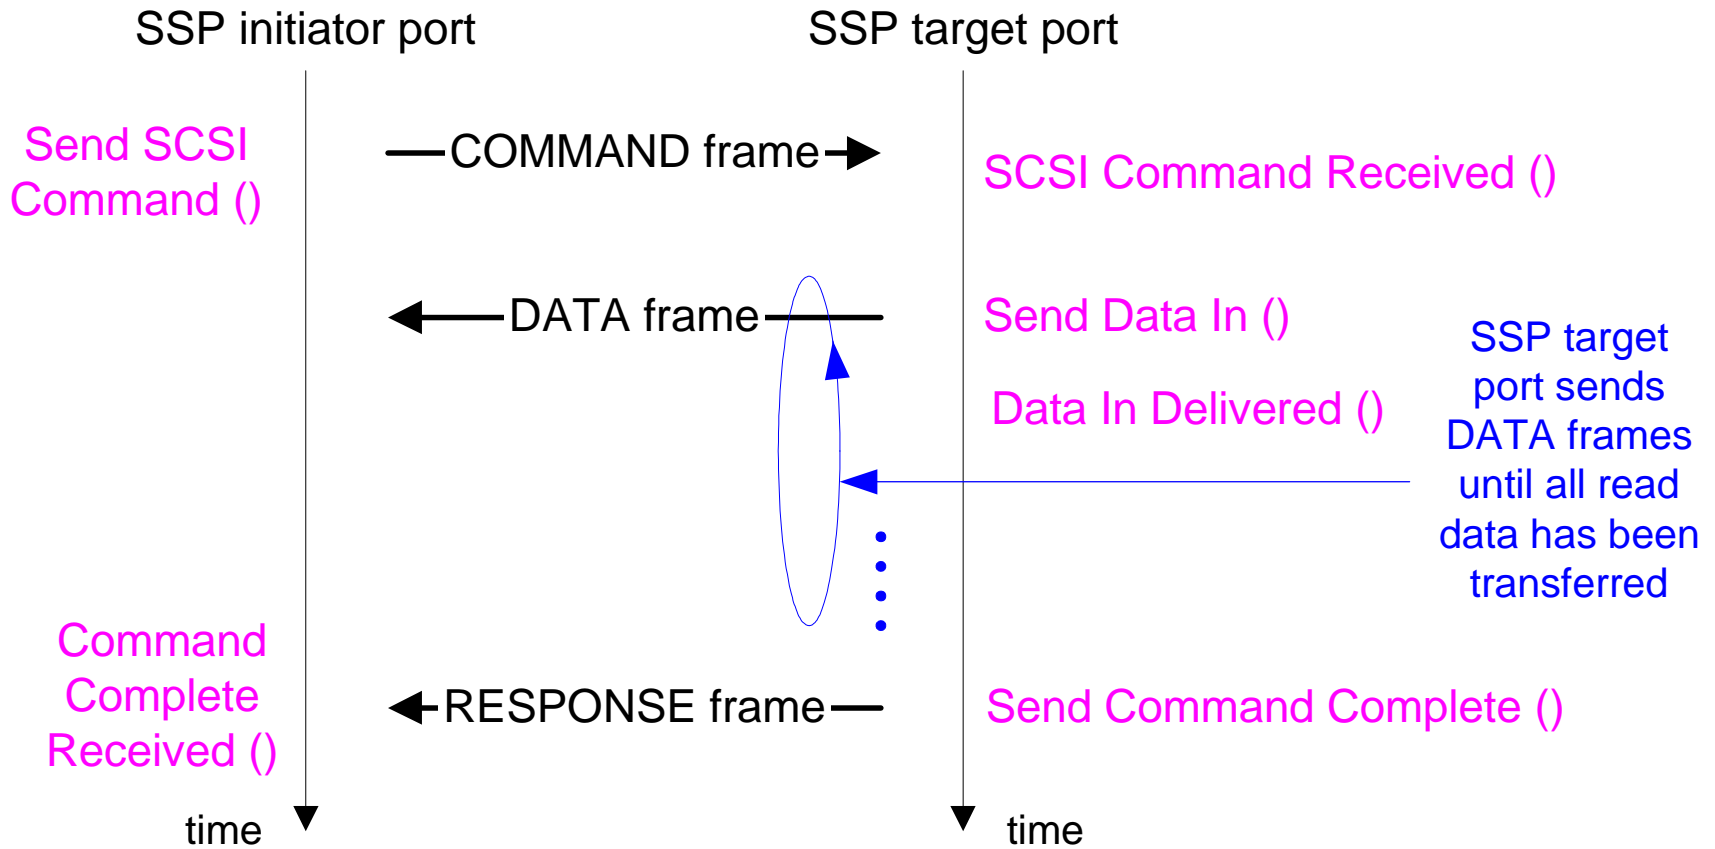


# SSP write command sequence

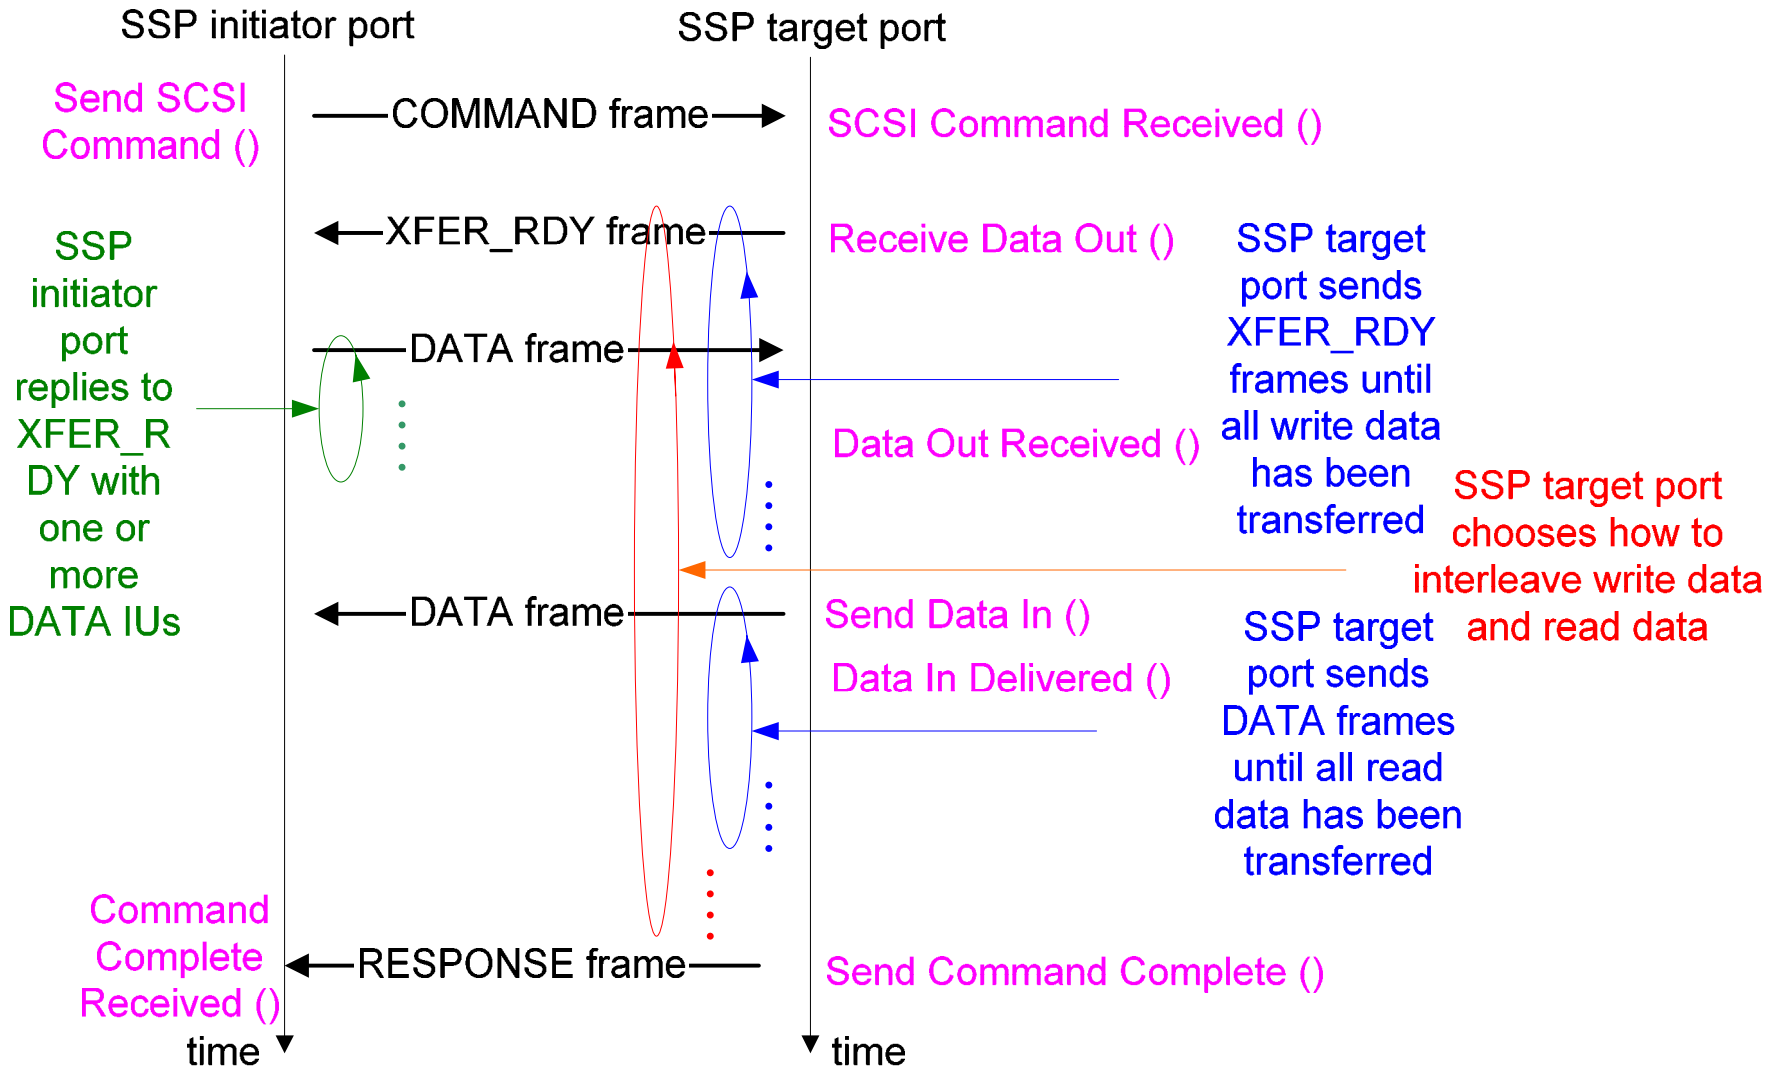




# SSP read command sequence



# SSP bidirectional command sequence





# 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>



i n v e n t