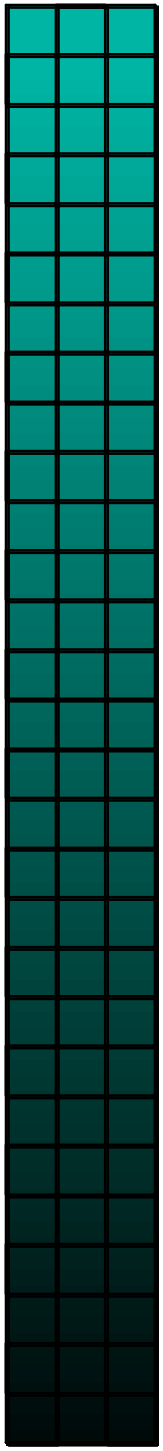

PCI to SSA Adapter

A single chip interface between PCI and Serial Storage Architecture

Adge Hawes, IBM Storage Subsystems

April 1996

SSA Overview

- Serial Link for Disk Drives
  - 20MB/s full duplex link, 40MB/s new version,
  - 160MB/s throughput for dual-port
  - 4 wire/port (Rx and Tx, differential), 8B/10B code
  - 25 metres/link, fibre-optic extension to 2.4Km
- String, Loop and Switch topology
- 128 byte frames, 8 byte overhead (6%)
  - Special characters for control
- Spatial Reuse for maximum performance
  - SAT Algorithm for fairness and arbitration
- ANSI Standardization, now shipping (IBM 7133).

SSA Shared Storage Array

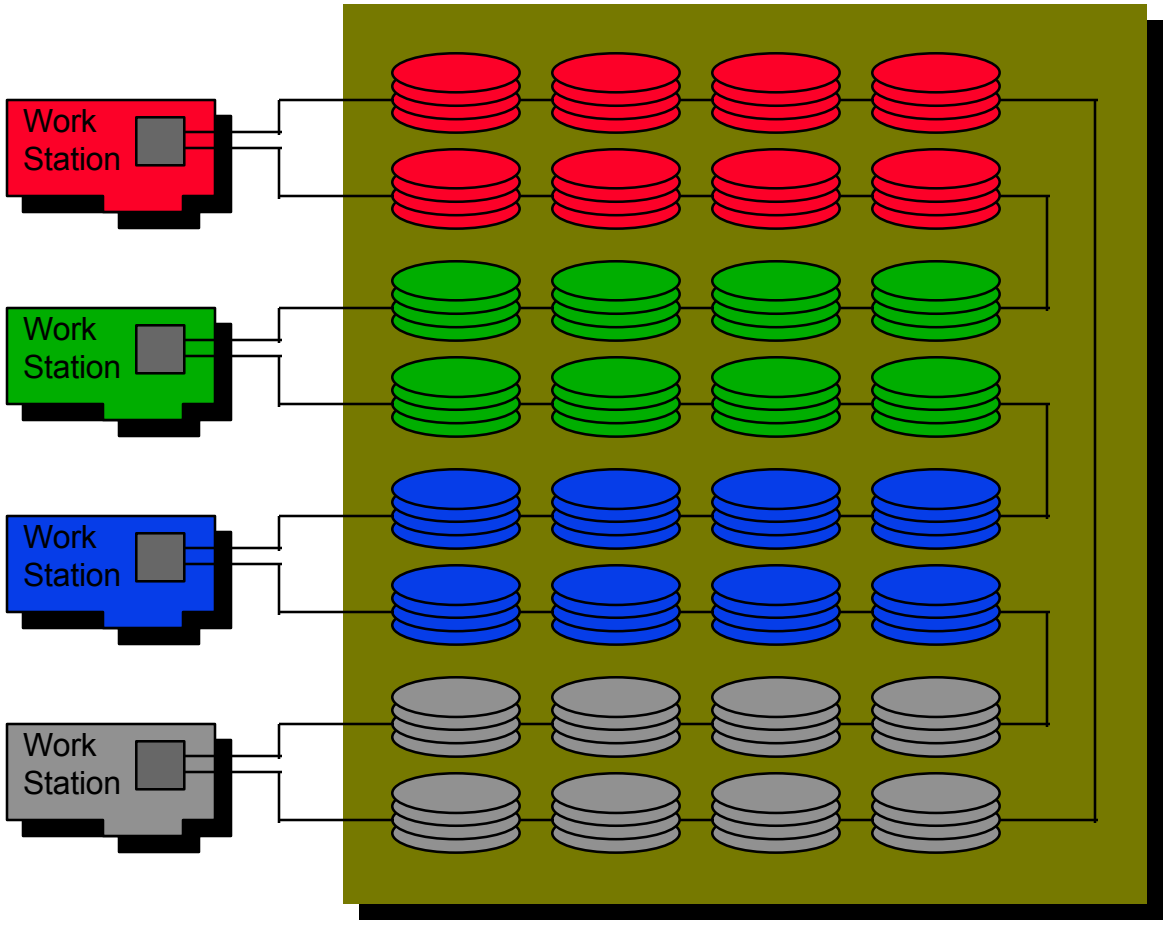

Spatial Reuse

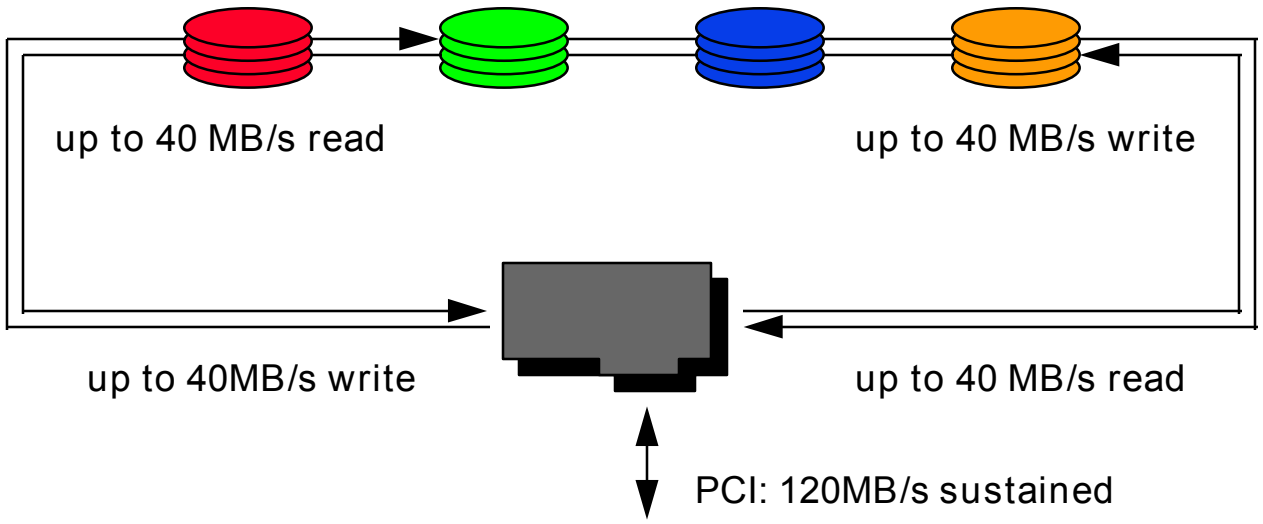

- Bandwidth available = 4x link bandwidth
- Limitation of PCI and memory

XORing disks and RAID

RAID 5 write without XORing disks or Spatial Reuse:

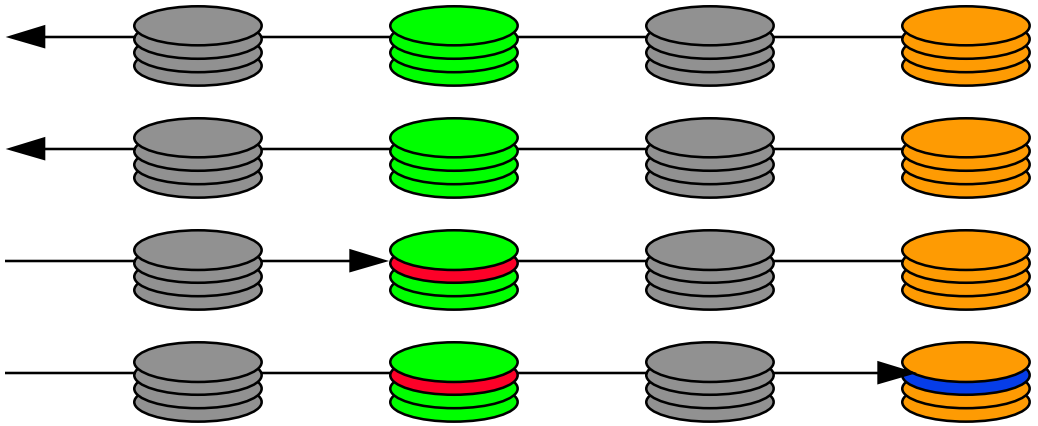

1. Read old Data (OD)
2. Read old Parity (OP)
3. Write new data (ND)
4. Write new Parity = ND ⊕ OD ⊕ OP

RAID 5 write with XORing disks and Spatial Reuse:

XOR new with old
Write new Data
Send result
XOR new Parity

SSA Traffic Flow

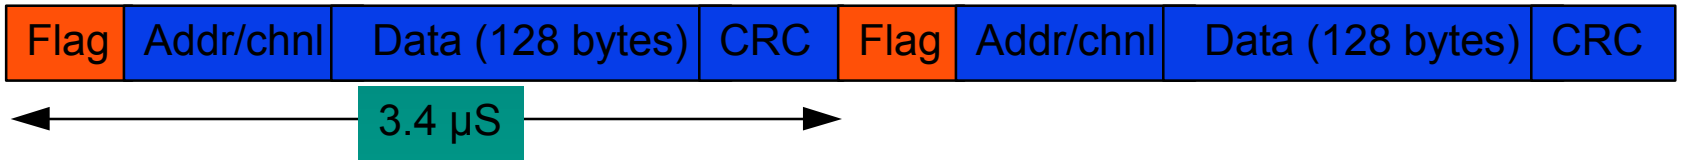

- All frames go through frame buffers
- “Through” frames - decrement address and pass on
- Address=0: Data for this node
- Inbound data for this adapter - must be captured and passed to the correct DMA channel on PCI
- Outbound data from this adapter - must be fetched from the correct DMA channel in PCI memory and inserted into output stream.

“Laredo”

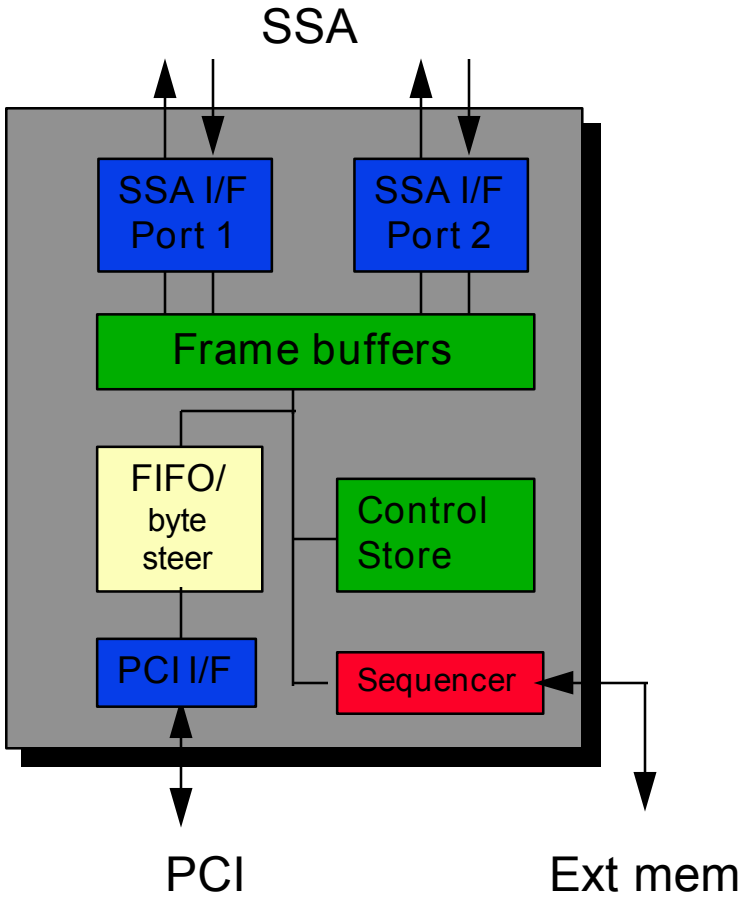

- 40/20 MB/s Dual-Port SSA
- PCI 2.1 compliant
- 64-bit Addressing
- 15 DMA Channels/port
- Internal Sequencer
  - 50MHz, 32-bit
  - 448x32 RAM
- External ROM/RAM
- 3.3v CMOS.

SSA/PCI speed matching

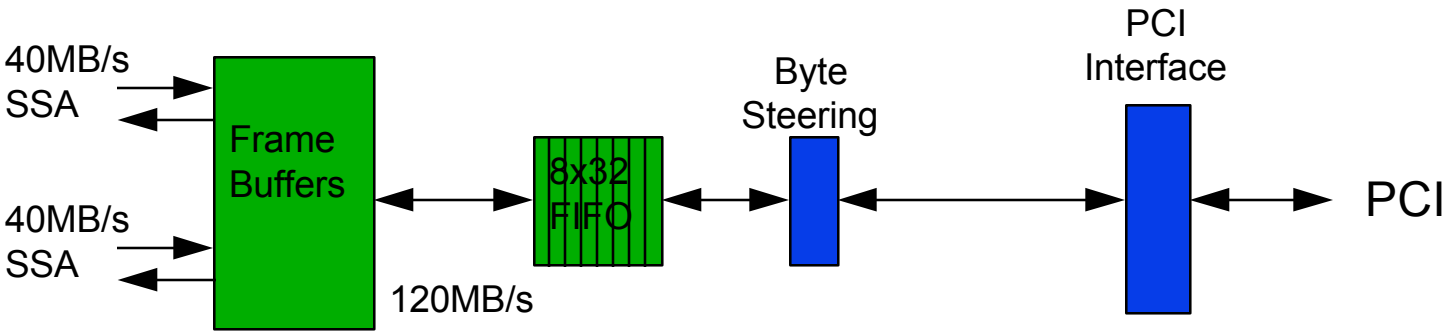

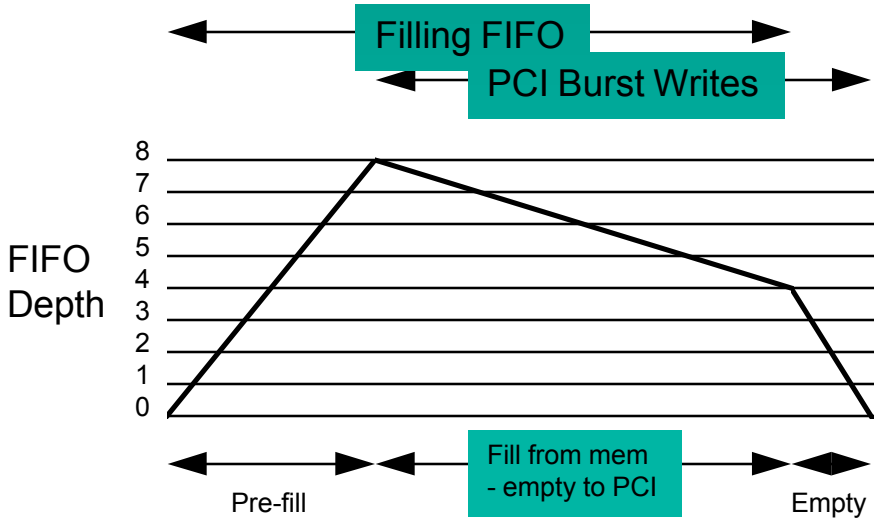

- 120MB/s memory bandwidth available
- FIFO allows 132MB/s burst for 128 bytes
- Byte steering allows unaligned transfers
- No wait states.

DMA

- A Channel Descriptor Block is fetched for each frame, in or out
- Inbound frames:
  - Channel component directs DMA engine
  - 128 bytes are posted into PCI memory

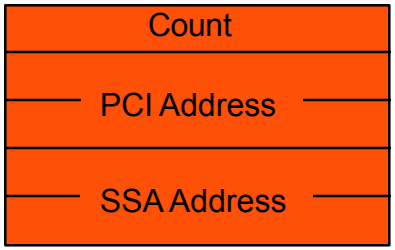

- Outbound frames:
  - Arbitration/Fairness selects channel to be serviced
  - More than one frame may be transmitted for channel
  - SAT Algorithm governs origination
- Scatter-gather boundaries handled by Sequencer
- Control Blocks - access to PCI by Sequencer.

PCI Communications

- Sequencer provides configuration address space
- External ROM/RAM for extra configuration and PCI BIOS
- Transactions via chained tag blocks in PCI memory

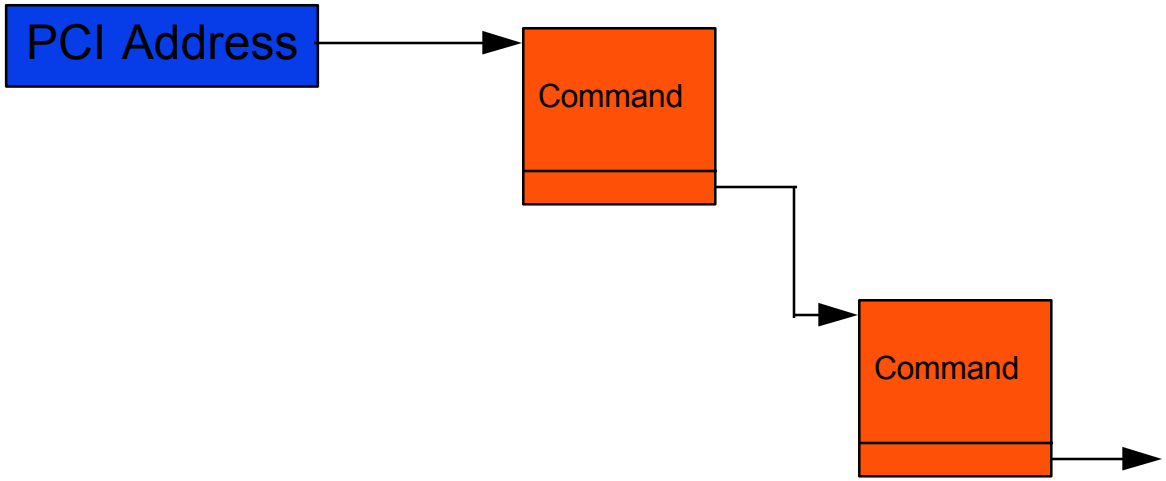

Conclusion

- Single-chip PCI-SSA Interface, PCI 2.1 compliant
- 40 and 20MB/s link operation
- Programmable via on-chip Sequencer
- Scatter-gather support for data transfer
- Good PCI citizen - no wait states
- External Memory Interface
- Low cost CMOS.