Disk Controllers

- HBA
- Based on disk interface

Disk Interfaces

- Hard disk drives are accessed over one of a number of bus types, including parallel **ATA (PATA, also called IDE or EIDE),** Serial ATA (SATA), SCSI, Serial Attached SCSI (SAS), and Fibre Channel. Bridge circuitry is sometimes used to connect hard disk drives to buses that they cannot communicate with natively, such as IEEE 1394 and USB.
- Back in the days of the ST-506 interface, the data encoding scheme was also important. The first ST-506 disks used Modified Frequency Modulation (MFM) encoding, and transferred data at a rate of 5 megabits per second. Later on, controllers using 2,7 RLL (or just "RLL") encoding increased the transfer rate by 50%, to 7.5 megabits per second; this also increased disk capacity by fifty percent.
- Many ST-506 interface disk drives were only specified by the manufacturer to run at the lower MFM data rate, while other models (usually more expensive versions of the same basic disk drive) were specified to run at the higher RLL data rate. In some cases, a disk drive had sufficient margin to allow the MFM specified model to run at the faster RLL data rate; however, this was often unreliable and was not recommended. (An RLL-certified disk drive could run on a MFM controller, but with 1/3 less data capacity and speed.)
- **Enhanced Small Disk Interface (ESDI)** also supported multiple data rates (ESDI disks always used 2,7 RLL, but at 10, 15 or 20 megabits per second), but this was usually negotiated automatically by the disk drive and controller; most of the time, however, 15 or 20 megabit ESDI disk drives weren't downward compatible (i.e. a 15 or 20 megabit disk drive wouldn't run on a 10 megabit controller). ESDI disk drives typically also had jumpers to set the number of sectors per track and (in some cases) sector size.

Disk Intrefaces

TYPES OF HARD DISKS			
Type of Interface	Encoding Method**	Transfer Rate (Per sec)	Range of Capacities
SATA (IDE)	RLL	150-300MB	40GB-1.2TB
PATA (IDE)	RLL	3-133MB	500MB-400GB
SCSI	RLL	5-320MB	20MB-300GB
Older Interfaces			
IPI	RLL	10-25MB	200MB-3GB
ESDI	RLL	1-3MB	80MB-2GB
SMD	RLL	1-4MB	200MB-2GB
IDE	RLL	1-8MB	40MB-1GB
ST506 RLL	RLL	937KB	30MB-200MB
ST506	MFM	625KB	5MB-100MB

Disk Interfaces

Acronym or abbreviation	Meaning	Description
SASI	Shugart Associates System Interface	Historical predecessor to SCSI.
SCSI	Small Computer System Interface	Bus oriented that handles concurrent operations.
SAS	Serial Attached SCSI	Improvement of SCSI, uses serial communication instead of parallel.
ST-506		Historical Seagate interface.
ST-412		Historical Seagate interface (minor improvement over ST-506).
ESDI	Enhanced Small Disk Interface	Historical; backwards compatible with ST-412/506, but faster and more integrated.
ATA	Advanced Technology Attachment	Successor to ST-412/506/ESDI by integrating the disk controller completely onto the device. Incapable of concurrent operations.
SATA	Serial ATA	Modification of ATA, uses serial communication instead of parallel.

Disk Interfaces

- Modern hard drives present a consistent interface to the rest of the computer, no matter what data encoding scheme is used internally. Typically a **DSP** in the electronics inside the hard drive takes the raw analog voltages from the read head and uses PRML and Reed–Solomon error correction to decode the sector boundaries and sector data, then sends that data out the standard interface. That DSP also watches the error rate detected by error detection and correction, and performs bad sector remapping, data collection for Self-Monitoring, Analysis, and Reporting Technology, and other internal tasks.

- SCSI originally had just one speed, 5 MHz (for a maximum data rate of five megabytes per second), but later this was increased dramatically. The SCSI bus speed had no bearing on the disk's internal speed because of buffering between the SCSI bus and the disk drive's internal data bus; however, many early disk drives had very small buffers, and thus had to be reformatted to a different interleave (just like ST-506 disks) when used on slow computers, such as early IBM PC compatibles and early Apple Macintoshes.

Disk Interfaces

- **ATA disks** have typically had no problems with interleave or data rate, due to their controller design, but many early models were incompatible with each other and couldn't run in a master/slave setup (two disks on the same cable). This was mostly remedied by the mid-1990s, when ATA's specification was standardised and the details began to be cleaned up, but still causes problems occasionally (especially with CD-ROM and DVD-ROM disks, and when mixing Ultra DMA and non-UDMA devices).

- **Serial ATA** does away with master/slave setups entirely, placing each disk on its own channel (with its own set of I/O ports) instead.

- **FireWire/IEEE 1394 and USB(1.0/2.0)** HDDs are external units containing generally ATA or SCSI disks with ports on the back allowing very simple and effective expansion and mobility. Most FireWire/IEEE 1394 models are able to **daisy-chain** in order to continue adding peripherals without requiring additional ports on the computer itself.

Disk interface families used in personal computers

- Notable families of disk interfaces include:
- **Historical bit serial interfaces** — connected to a hard disk drive controller with three cables, one for data, one for control and one for power. The HDD controller provided significant functions such as serial to parallel conversion, data separation and track formatting, and required matching to the drive in order to assure reliability.
  - **ST506** used MFM (Modified Frequency Modulation) for the data encoding method.
  - **ST412** was available in either MFM or RLL (Run Length Limited) variants.
  - **Enhanced Small Disk Interface (ESDI)** was an interface developed by **Maxtor** to allow faster communication between the PC and the disk than MFM or RLL.

Modern bit serial interfaces

- Modern **bit serial interfaces** — connect to a host bus adapter (today typically integrated into the "south bridge") with two cables, one for data/control and one for power.
  - **Fibre Channel (FC)**, is a successor to parallel SCSI interface on enterprise market. It is a serial protocol. In disk drives usually the Fibre Channel Arbitrated Loop (FC-AL) connection topology is used. FC has much broader usage than mere disk interfaces, it is the cornerstone of storage area networks (SANs). Recently other protocols for this field, like iSCSI and ATA over Ethernet have been developed as well. Confusingly, drives usually use *copper* twisted-pair cables for Fibre Channel, not fibre optics. The latter are traditionally reserved for larger devices, such as servers or disk array controllers.
  - **Serial ATA (SATA).** The SATA data cable has one data pair for differential transmission of data to the device, and one pair for differential receiving from the device, just like EIA-422. That requires that data be transmitted serially. The same differential signaling system is used in RS485, LocalTalk, USB, Firewire, and differential SCSI.
  - **Serial Attached SCSI (SAS**). The SAS is a new generation serial communication protocol for devices designed to allow for much higher speed data transfers and is compatible with SATA. SAS uses serial communication instead of the parallel method found in traditional SCSI devices but still uses SCSI commands.

Word serial interfaces (parallel)

- **Word serial interfaces** — connect to a host bus adapter (today typically integrated into the "south bridge") with two cables, one for data/control and one for power. The earliest versions of these interfaces typically had a 16 bit parallel data transfer to/from the drive and there are 8 and 32 bit variants. Modern versions have serial data transfer. The word nature of data transfer makes the design of a host bus adapter significantly simpler than that of the precursor HDD controller.
  - **Integrated Drive Electronics (IDE),** later renamed to ATA, and then later to PATA ("parallel ATA", to distinguish it from the new Serial ATA). The original name reflected the innovative integration of HDD controller with HDD itself, which was not found in earlier disks. Moving the HDD controller from the interface card to the disk drive helped to standardize interfaces, including reducing the cost and complexity. The 40 pin IDE/ATA connection of PATA transfers 16 bits of data at a time on the data cable. The data cable was originally 40 conductor, but later higher speed requirements for data transfer to and from the hard drive led to an "ultra DMA" mode, known as UDMA, which required an 80 conductor variant of the same cable; the other conductors provided the grounding necessary for enhanced high-speed signal quality. The interface for 80 conductor only has 39 pins, the missing pin acting as a key to prevent incorrect insertion of the connector to an incompatible socket, a common cause of disk and controller damage.
  - EIDE was an unofficial update (by Western Digital) to the original IDE standard, with the key improvement being the use of direct memory access (DMA) to transfer data between the disk and the computer without the involvement of the CPU, an improvement later adopted by the official ATA standards. By directly transferring data between memory and disk, DMA does not require the CPU/program/operating system to leave other tasks idle while the data transfer occurs.
  - **Small Computer System Interface (SCSI)**, originally named SASI for Shugart Associates System Interface, was an early competitor of ESDI. SCSI disks were standard on servers, workstations, and Apple Macintosh computers through the mid-90s, by which time most models had been transitioned to IDE (and later, SATA) family disks. Only in 2005 did the capacity of SCSI disks fall behind IDE disk technology, though the highest-performance disks are still available in SCSI and Fibre Channel only. The length limitations of the data cable allows for external SCSI devices. Originally SCSI data cables used single ended data transmission, but server class SCSI could use differential transmission, either low voltage differential (LVD) or high voltage differential (HVD).

Older Interface and controllers: MFM

- **MFM** encoding with serial transfer.
- Must set a Drive Select number.
- Required terminating resistors at both ends.
- 34 line control cable and 20 pin data cable.
- MFM: 17 sec/track 512 bytes/sec.

- **RLL**: 25 or 26 sec/track
- 7.5 Mbps transfer rate.
- Required a low level format.

- **Data cable**

Control cable

An MFM disk controller, an example of a 16-bit ISA / AT-bus card

MFM Controller WD1006

Older Interface and controllers: ESDI

- ESDI 24Mbps transfer rate.
- Some were compatible with ST-506.
- 34 line control cable and 20 pin data cable.
- St-506 / ESDI cable pinout

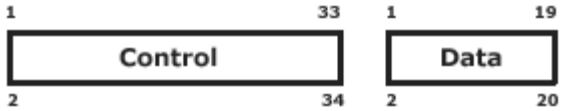

- St-506 / ESDI disk connections

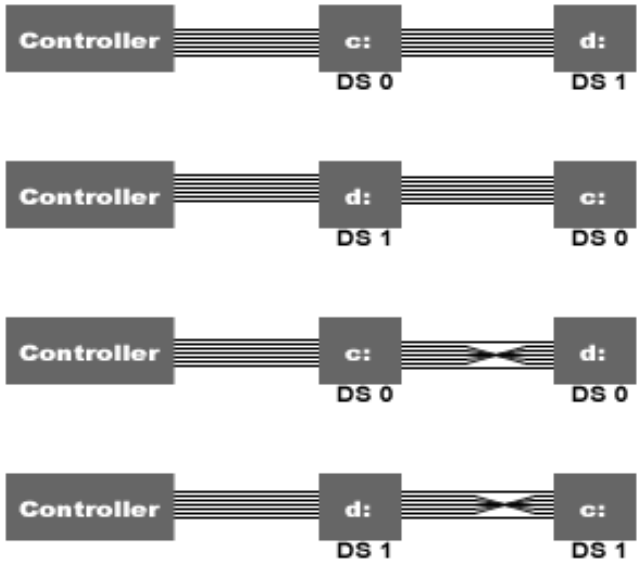

Koncepti arhitekture disk kontrolera

- 1. Compact controllers
- 2. Tracking-Buffer controllers (FIFO)
- 3. Full-Caching controllers

Compact controllers: Serial to Parallel: Idea for IDE

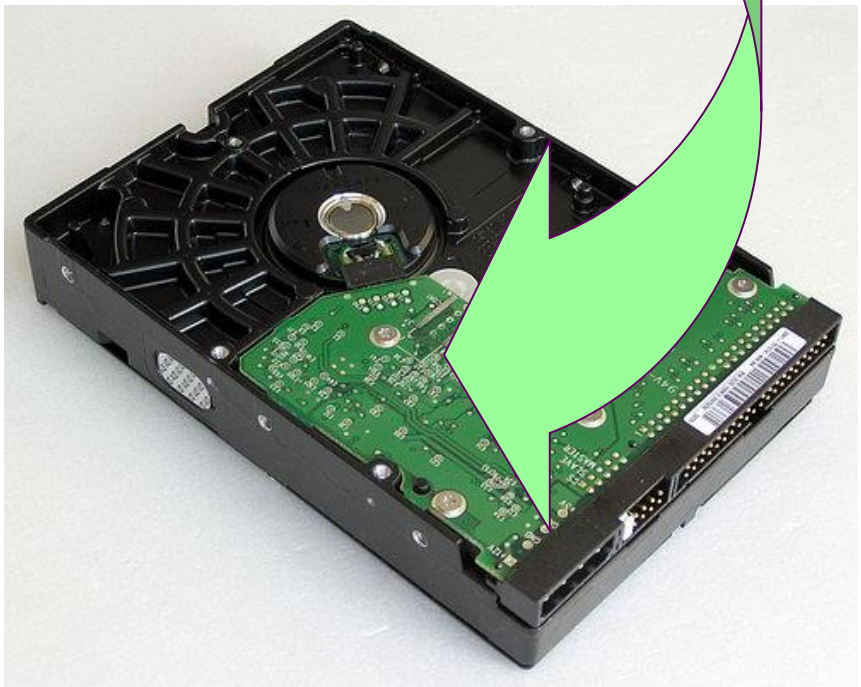

1. Compact controllers

- 1. Uprošćeni
- 2. Mali bafer (Nemaju veliki bafer)
- 3. Nemaju sopstveni CPU
- Karakteristike
  - mali bafer
  - prosta dekoderska logika
  - jednostavnost
- Primer
- ATA

System-board architecture for EIDE drives

Tipican predstavnik - ATA

ATA Controller - external

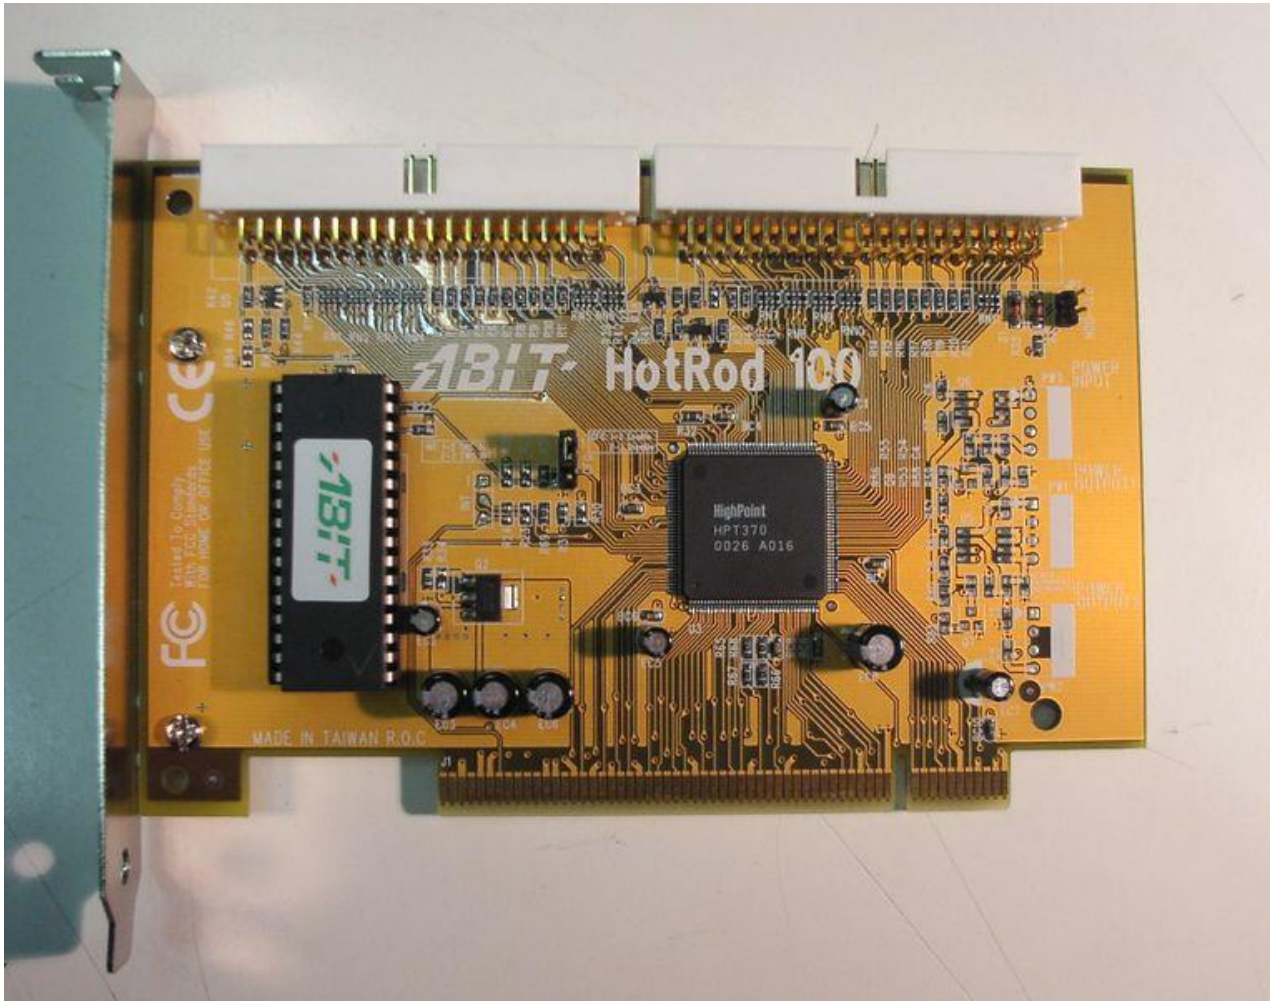

2. FIFO architecture

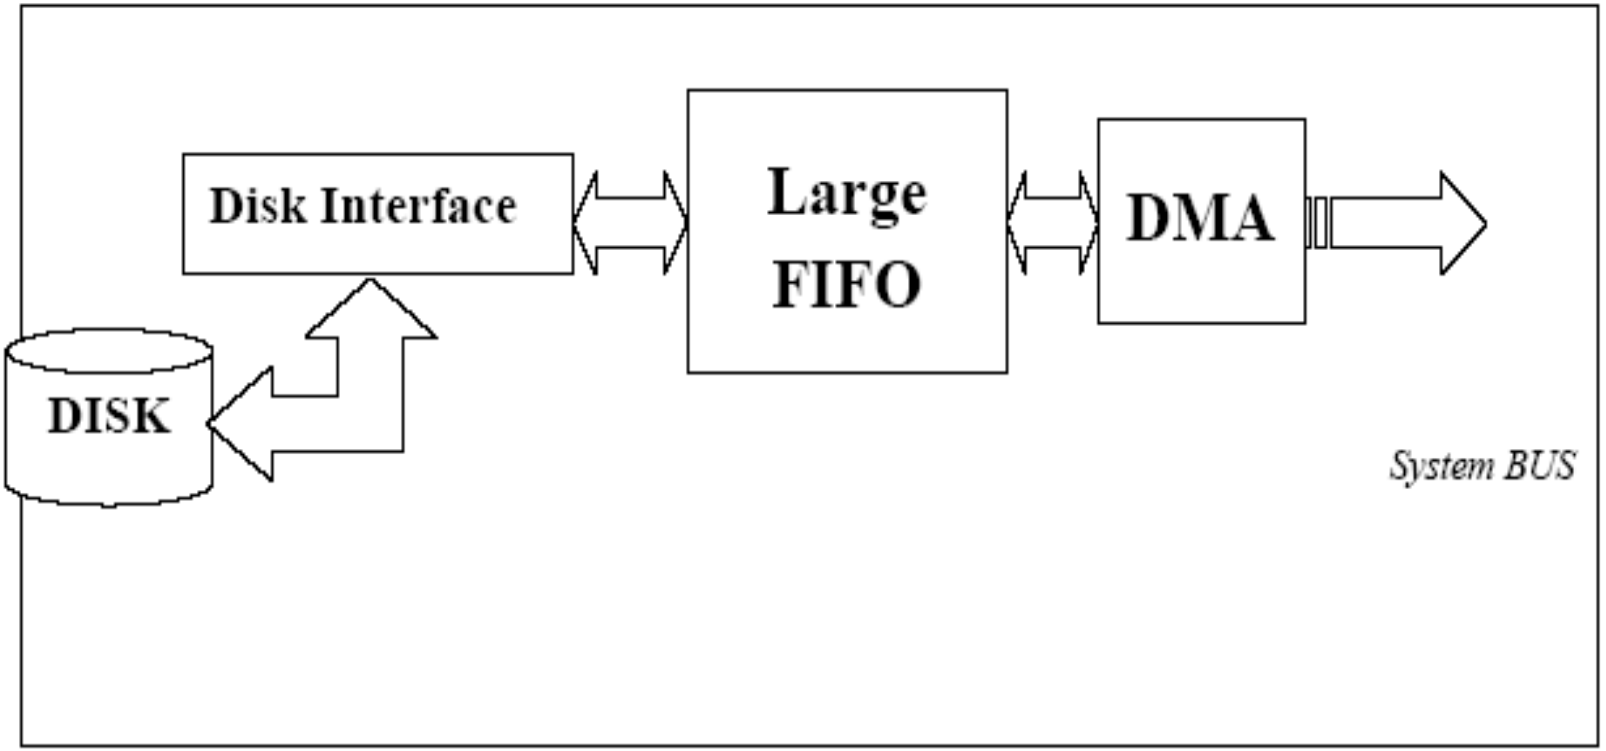

1. **Speed matching**
2. **Buffer size maching**
3. **Ideal for sequential performaces**

Kontroler scsi for isa bus

SCSI HBA, an example of a 32-bit PCI card

SCSI u160


ULTRA2-LVD/SE
FAST/ULTRA-SE
adaptec
04CVF
ASSEMBLED IN SINGAPORE

SCSI u320

3. Full-Caching Architecture

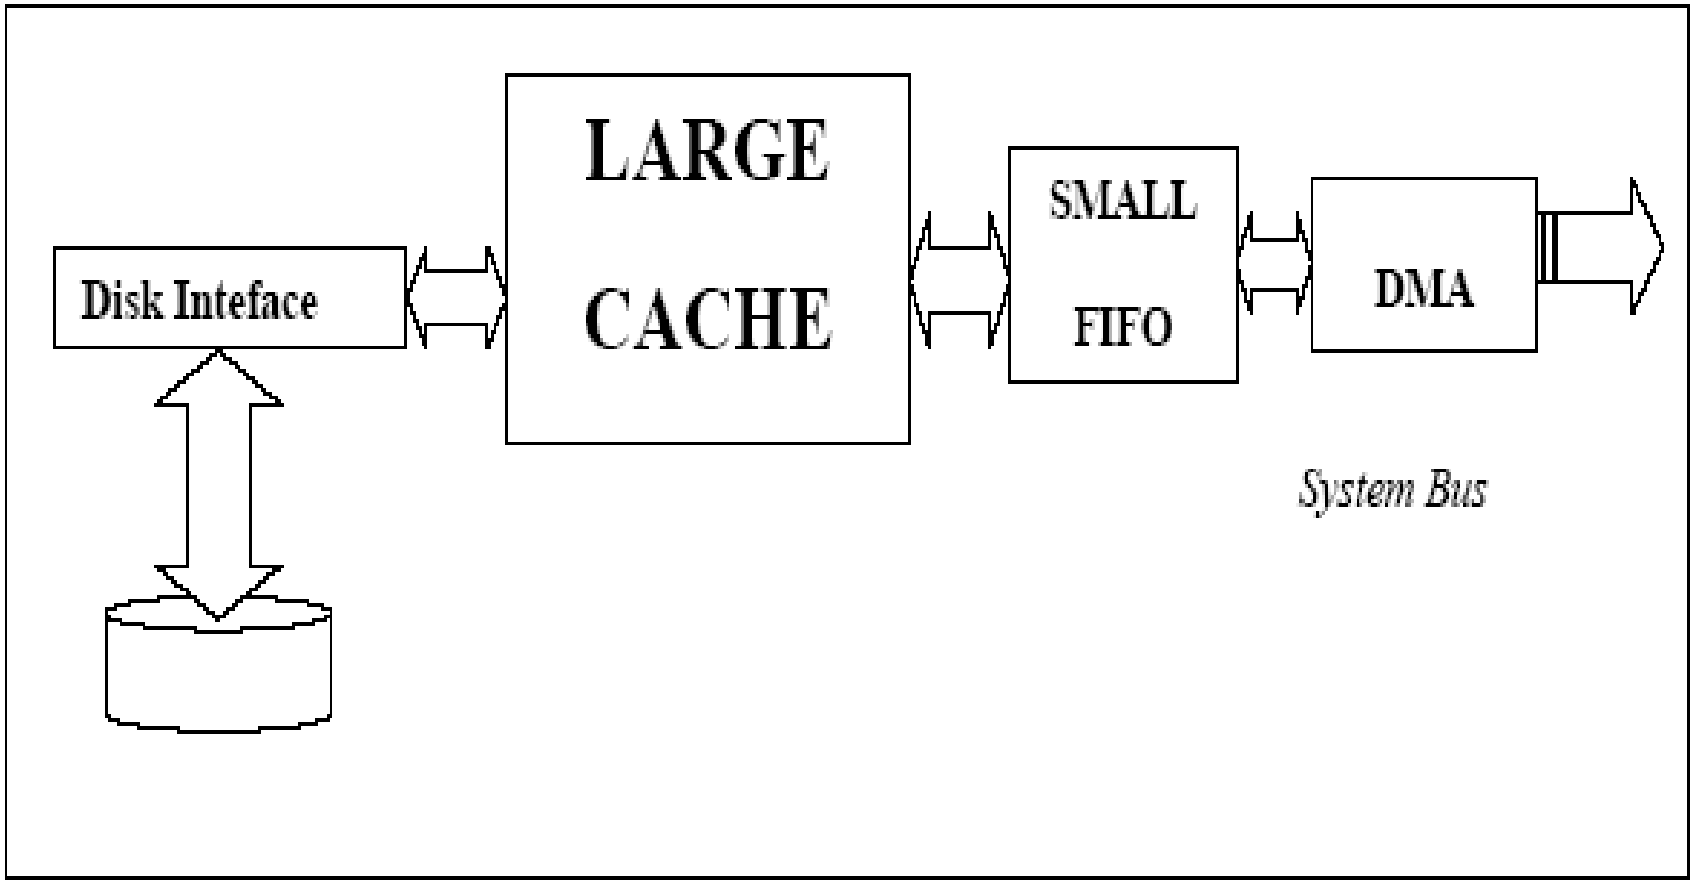

1. **Non-optimal architecture**
2. **Only as RAID caching**

3. Full-Caching Architecture

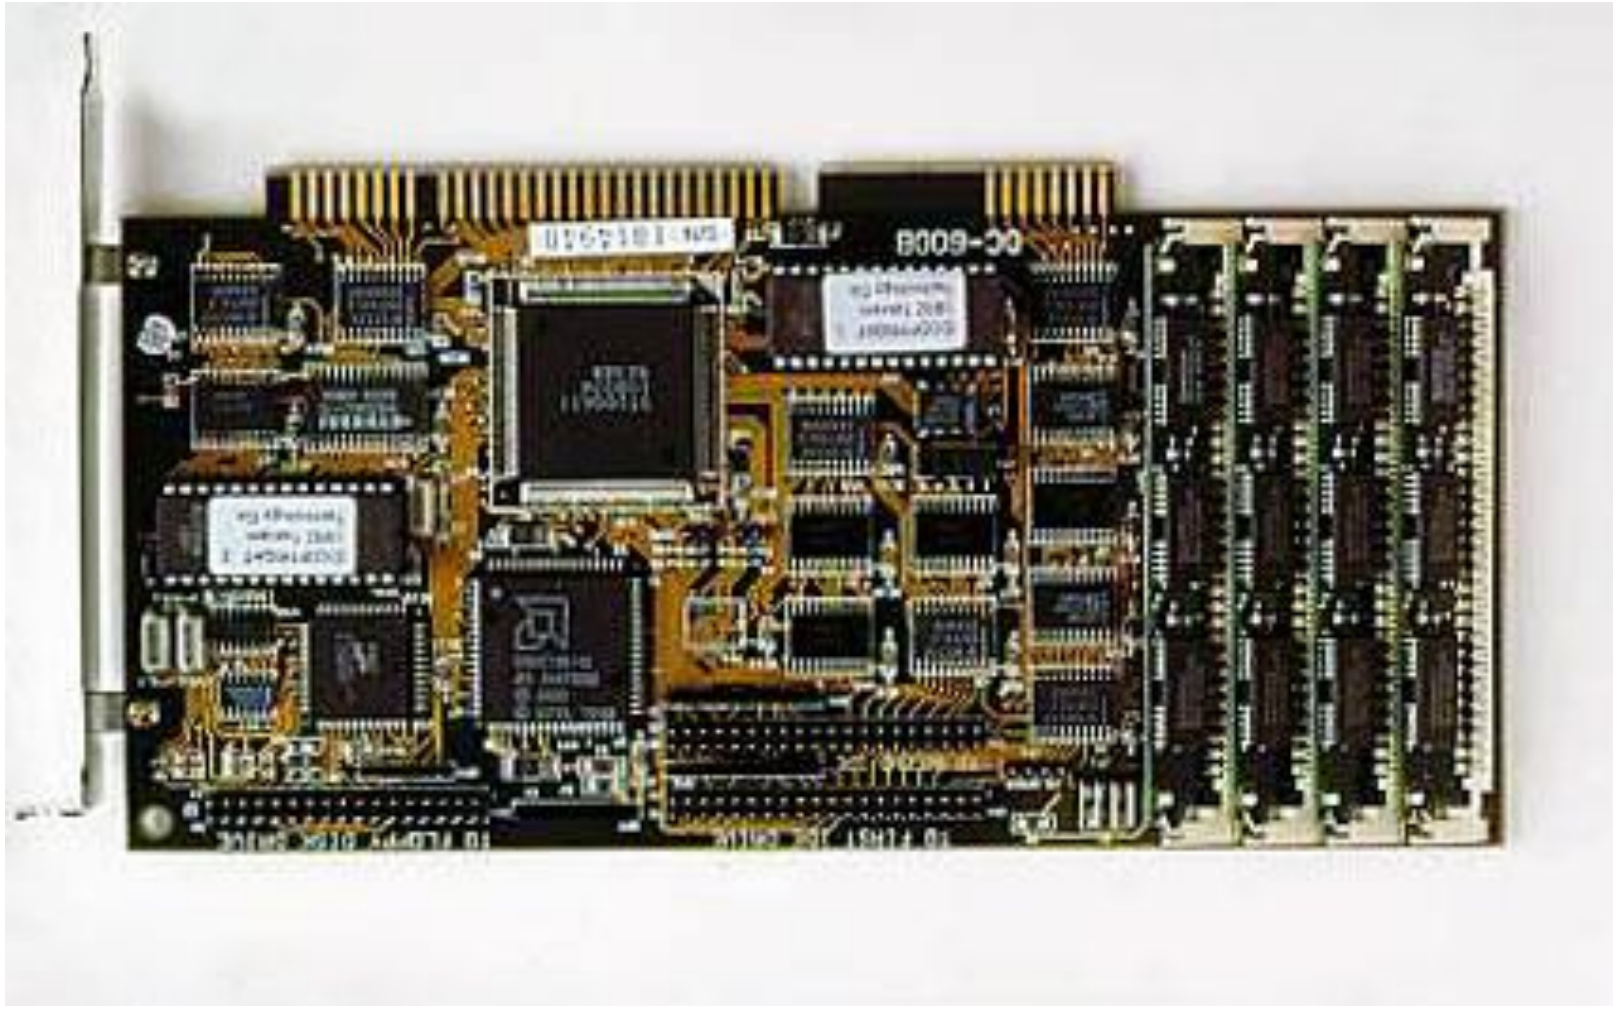

Disk caching

- Disk caches are usually just a part of main memory made up of common dynamic RAM (DRAM) chips, whereas memory caches (CPU caches) use higher-speed static RAM (SRAM) chips.

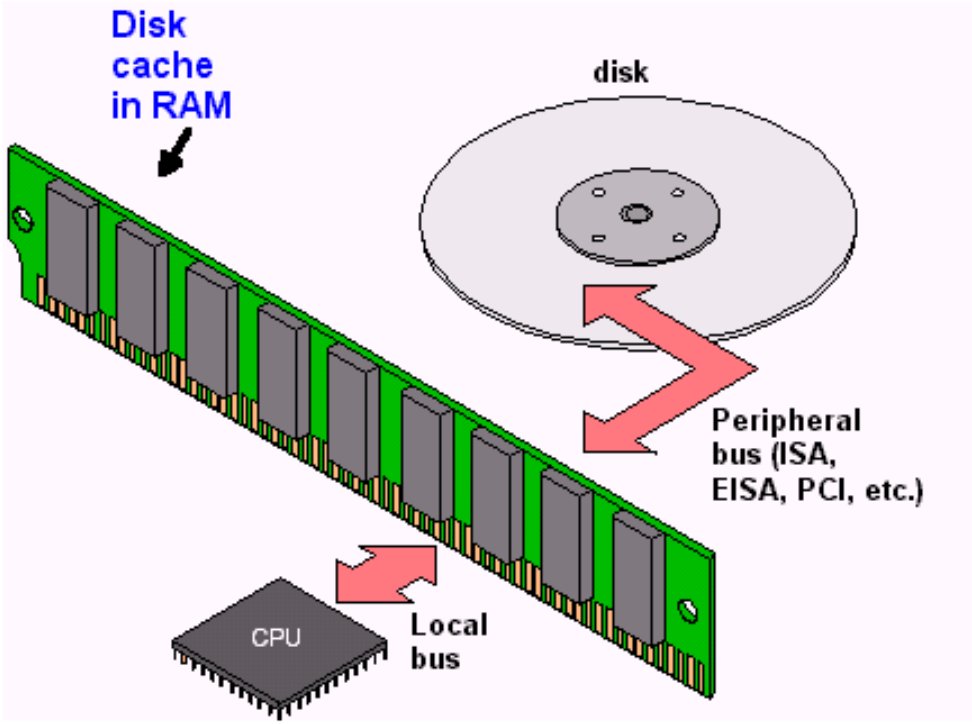

Full-Caching + RAID

Full-Caching + RAID

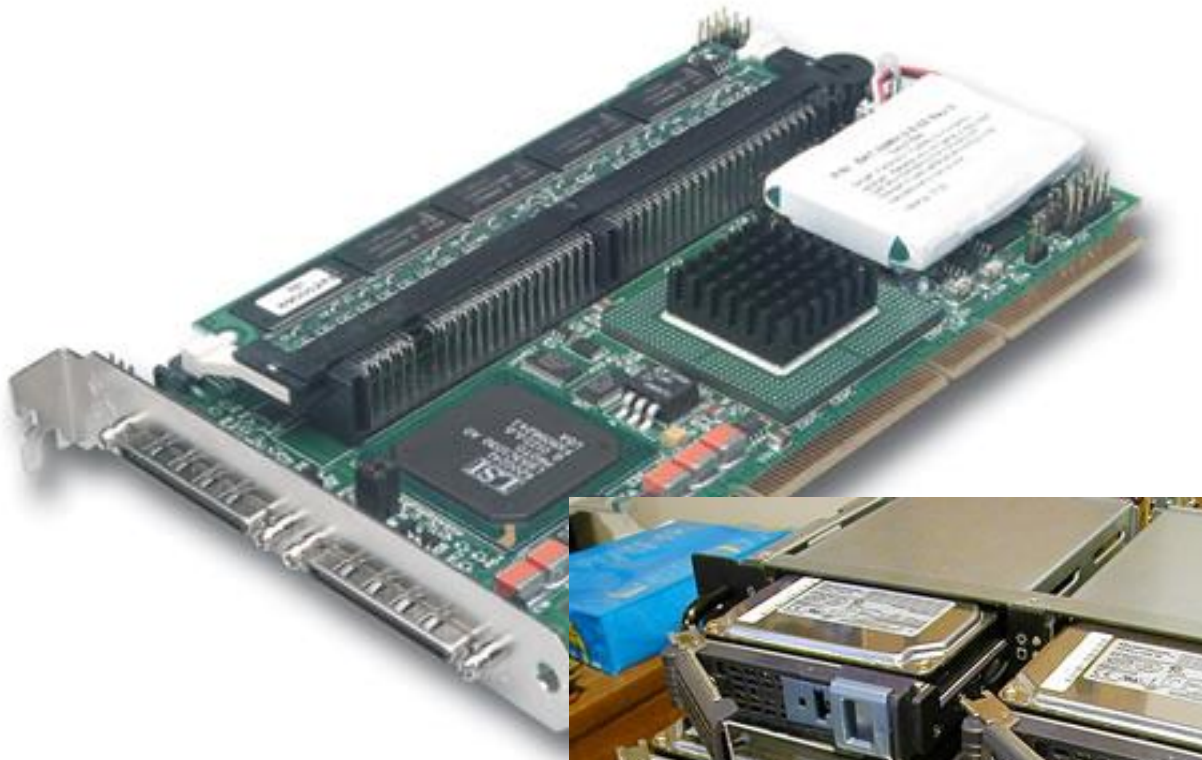

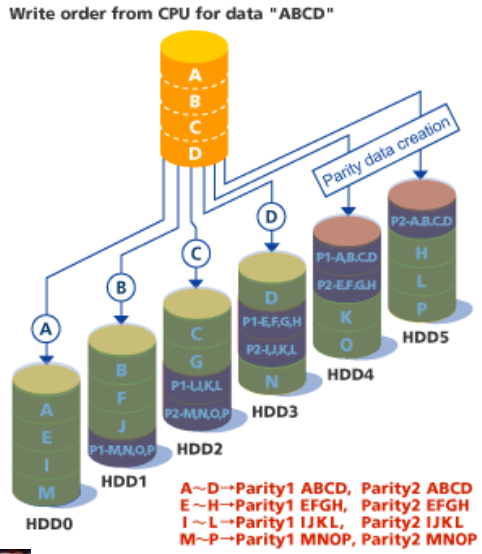

Full-Caching + RAID

Full-Caching + RAID

Full-Caching + RAID

Full-Caching + RAID

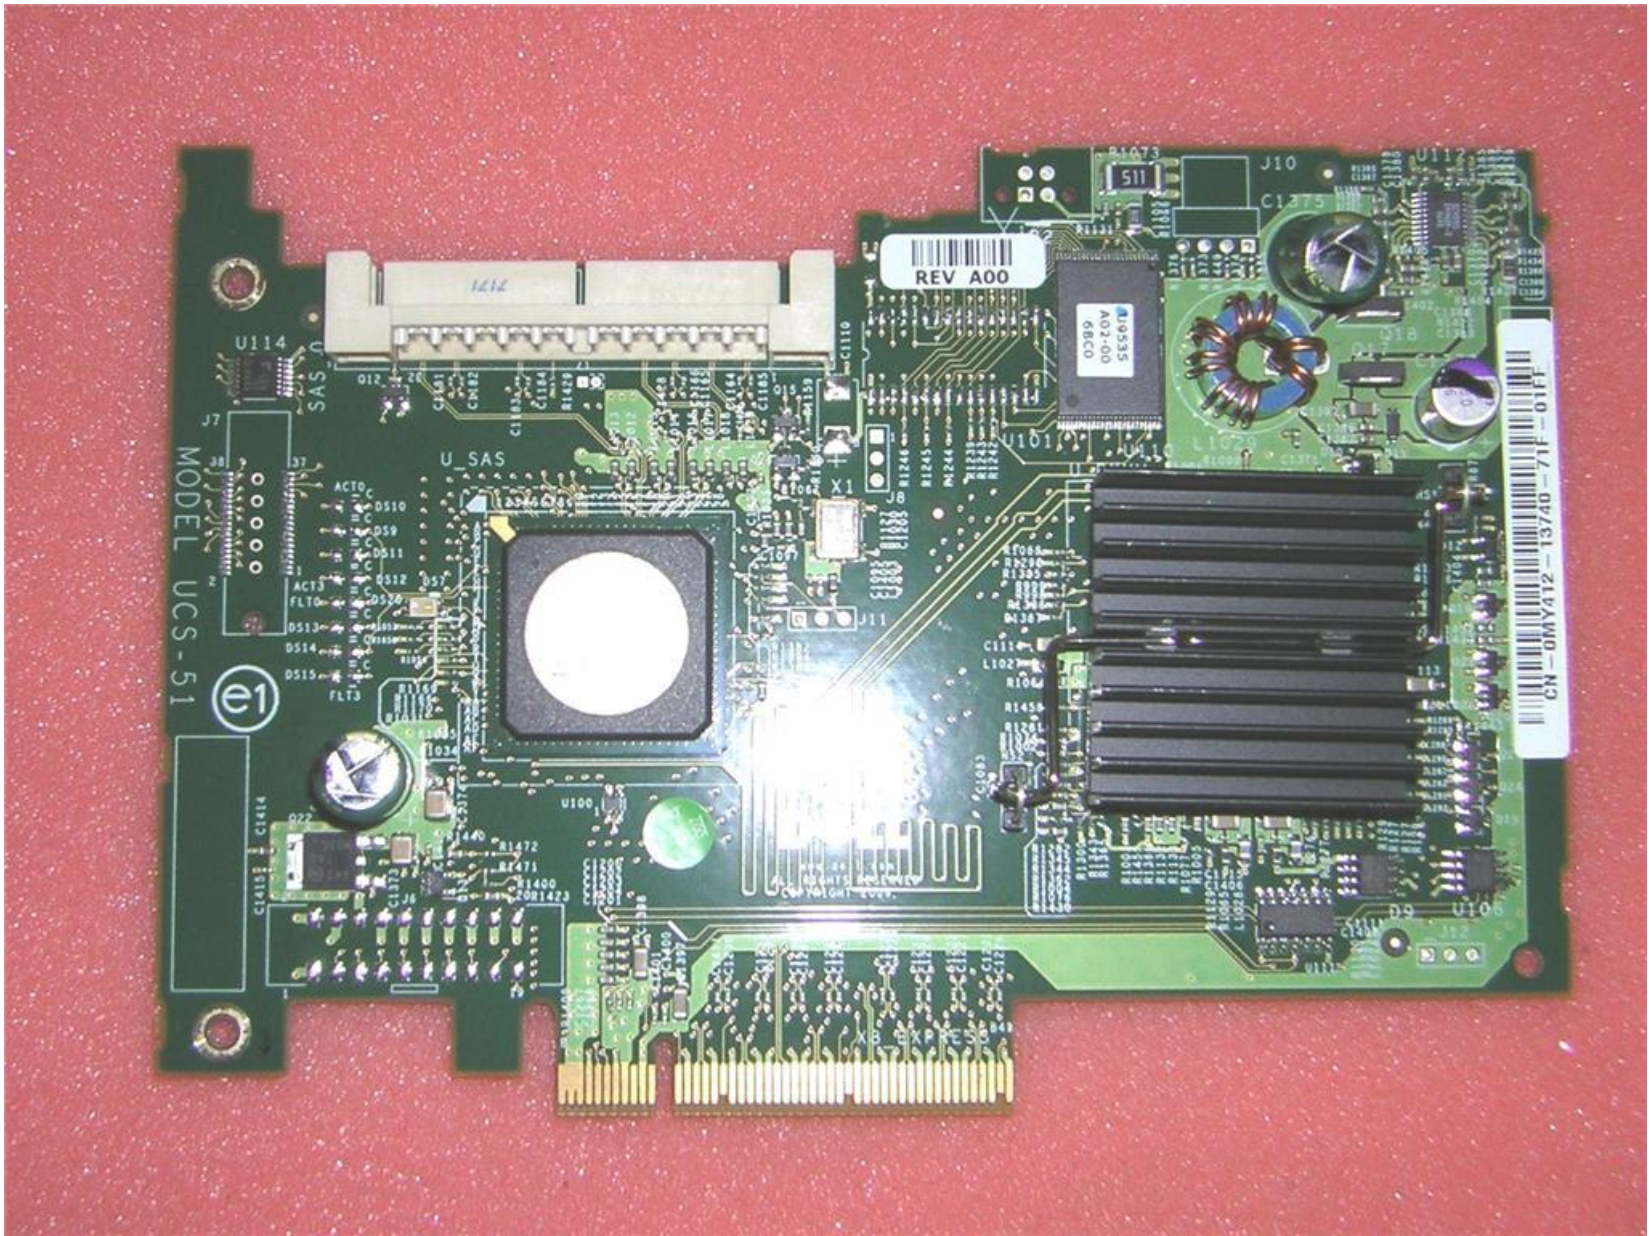

Full-Caching + RAID

Modern HBA function

- Speed matching
- Buffer size matching
- Command queuing
- Command reordering
- Hardware RAID capability

Back to serial

- **Key features**
  - High-speed 150MB/s (1.5Gbps)
  - RAID 0 / 1
  - supports larger than 137GB hard disk
  - external SATA port
  - PATA port available

Back to serial

- **Specifications**:
  - fully compliant with Serial ATA (SATA) 1.0 specification with data rates up to 150MB/s (1.5Gbps)
  - 2 independent SATA channels supporting one SATA device on each channel
  - ATA channel is shared with SATA channel : **1 SATA + 1 ATA channel when ATA is in use**
  - ·supports hard disk larger than 137GB
  - independent 256-byte FIFOs (32bit*64 deep) per Serial ATA channel for host reads and writes compliant with PCI specification 2.2 integrated PCI DMA engines
  - 32bit, 33 / 66 MHz fully compliant PCI host interface
  - RAID operating system drivers : Windows98 / Me / NT4.0 / 2000 / XP
  - RAID 0 (striping) / RAID 1 (mirroring) software included

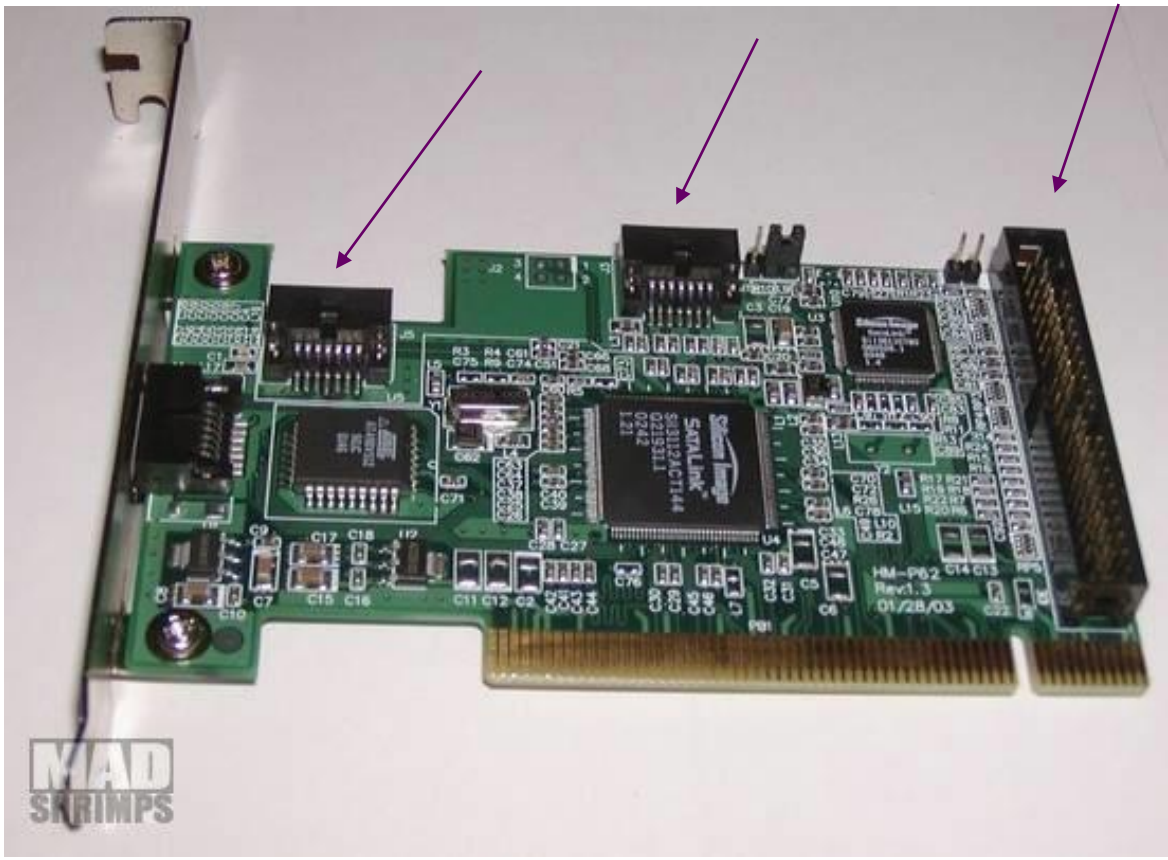

Back to serial

MiniPCI-SATA.jpg

FC controller (QLA 2200 Fibre Channel HBA)