

**Quantum®**

**Overlap IDE  
ATA Committee**

**June 2, 1995**

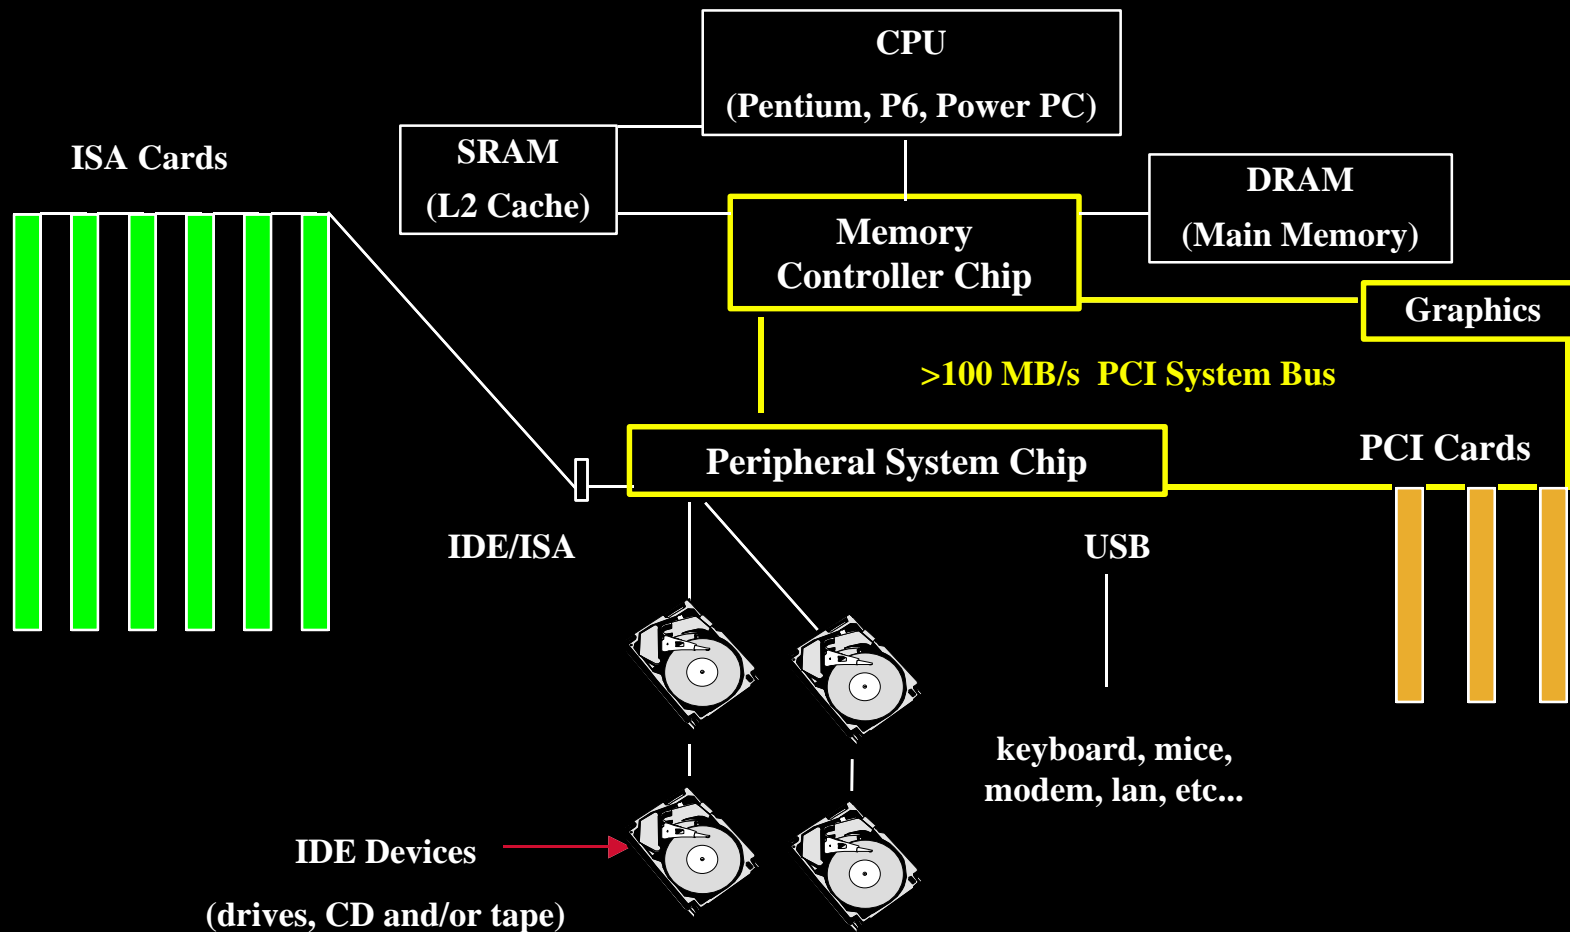
James McGrath

Systems Engineer  
Strategic and Technical Marketing

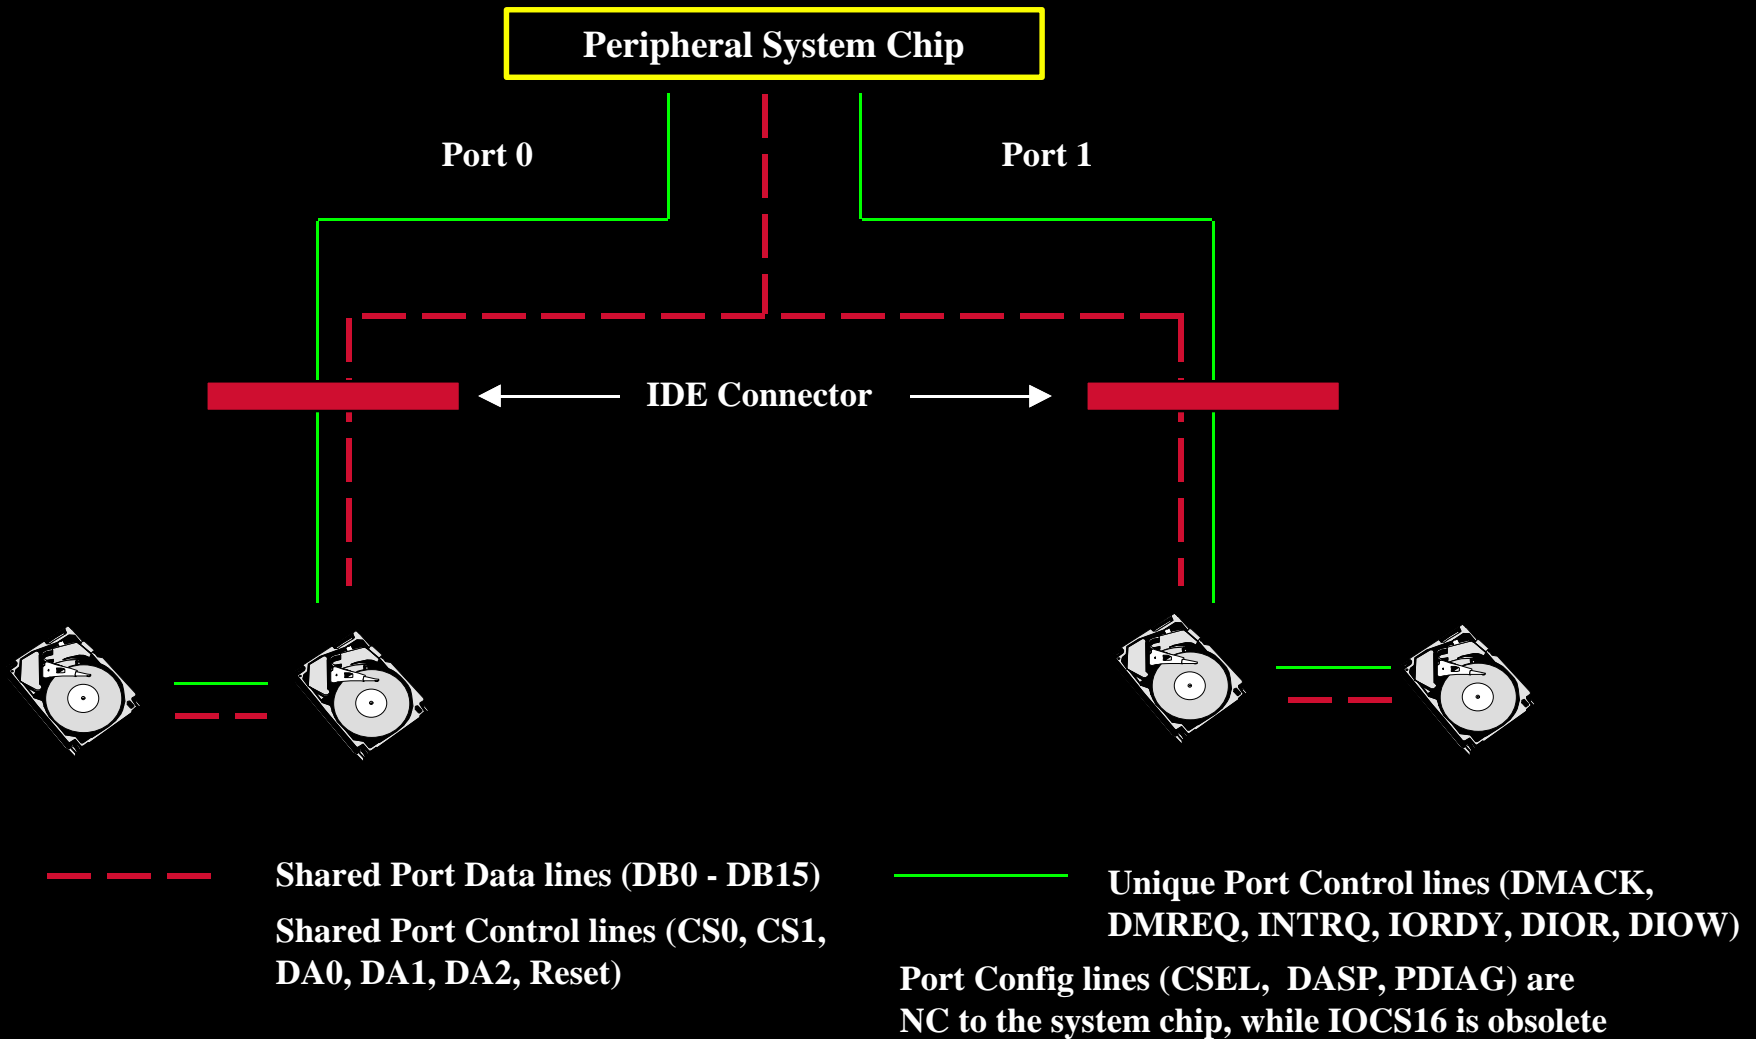
Quantum  
500 McCarthy Blvd  
Milpitas, CA 95035

phone: 408-894-4504  
fax: 408-894-4990  
internet: JMCGRATH@QNTM.COM

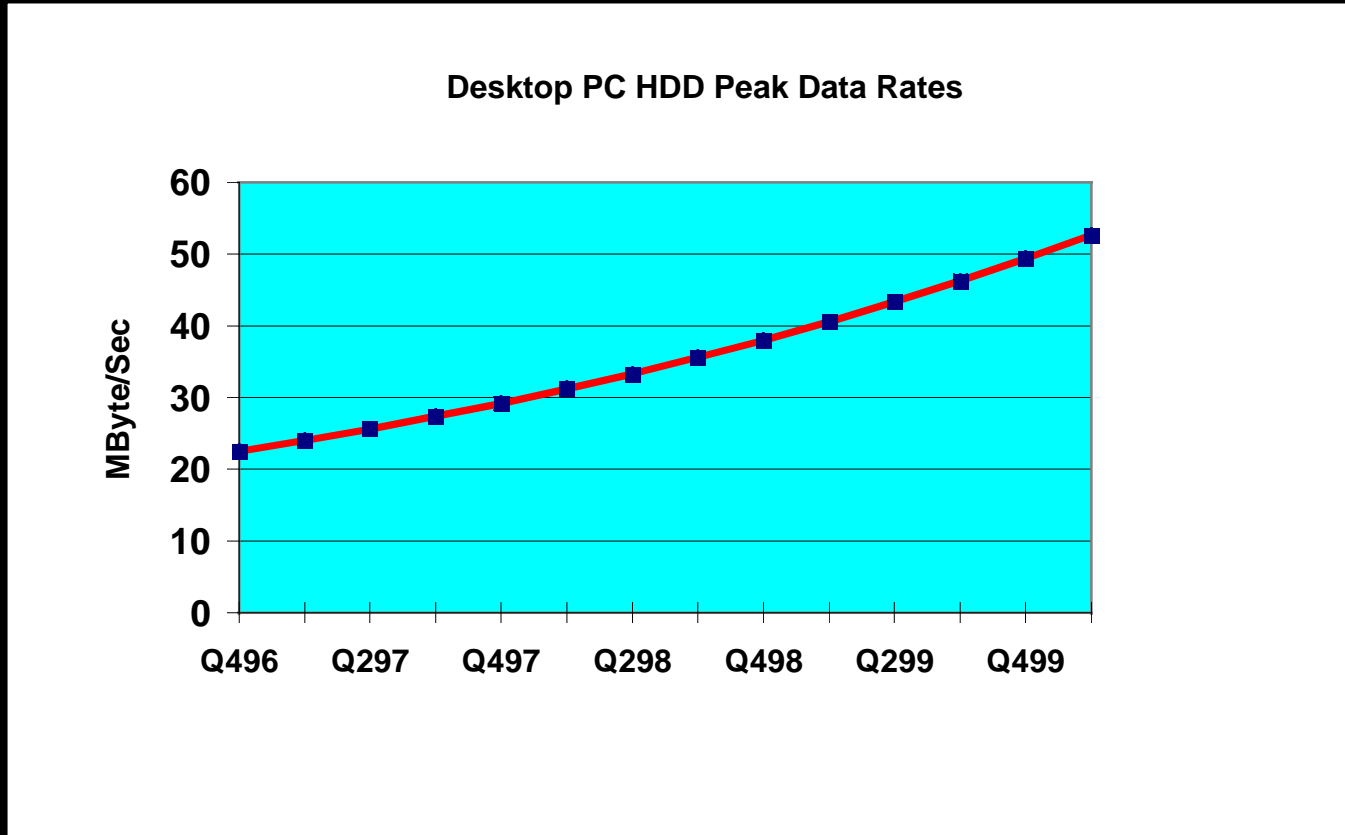
# PC Architecture - IDE Drive (Fast ATA/PCI Local Bus)



# PC Architecture - Focus on IDE Bus Today



# Read Channel Speed Evolution



## Read Channel Speed Evolution

- Assume density improvement is equally the responsibility of tpi and bpi increases over time.
- Implies that bpi, and thus read channel speed, increases at a 26% CAGR (doubling in 3 years).
- Net Host transfer rate must be > on track data transfer rate (about 75% to 80% of the read channel speed).

1995 - 8 MB/s

1996 - 12 MB/s

1997 - 16 MB/s

2000 - 32 MB/s


## Net Host Transfer Rate

- When accessing data sequentially, the net host transfer rate must be greater than the net disk transfer rate in order to avoid buffer overruns and slipping revolutions.
- Net host transfer rate factors in device and host overheads.

Device overhead has been reduced by automation.

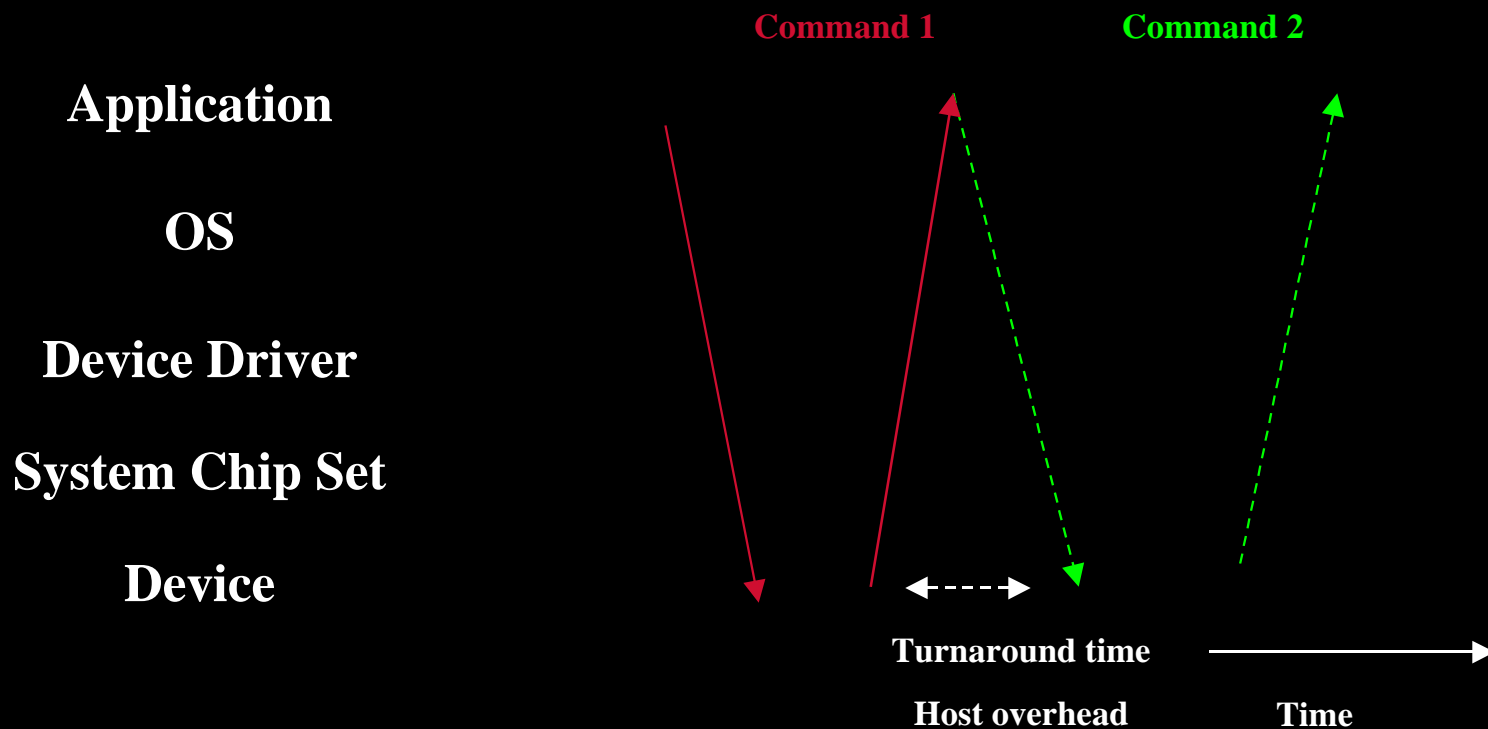
Host overhead between commands is still high since the OS is involved.

4K DMA read (prefetch hit) at 16 MB/s -> 7.6 MB/s



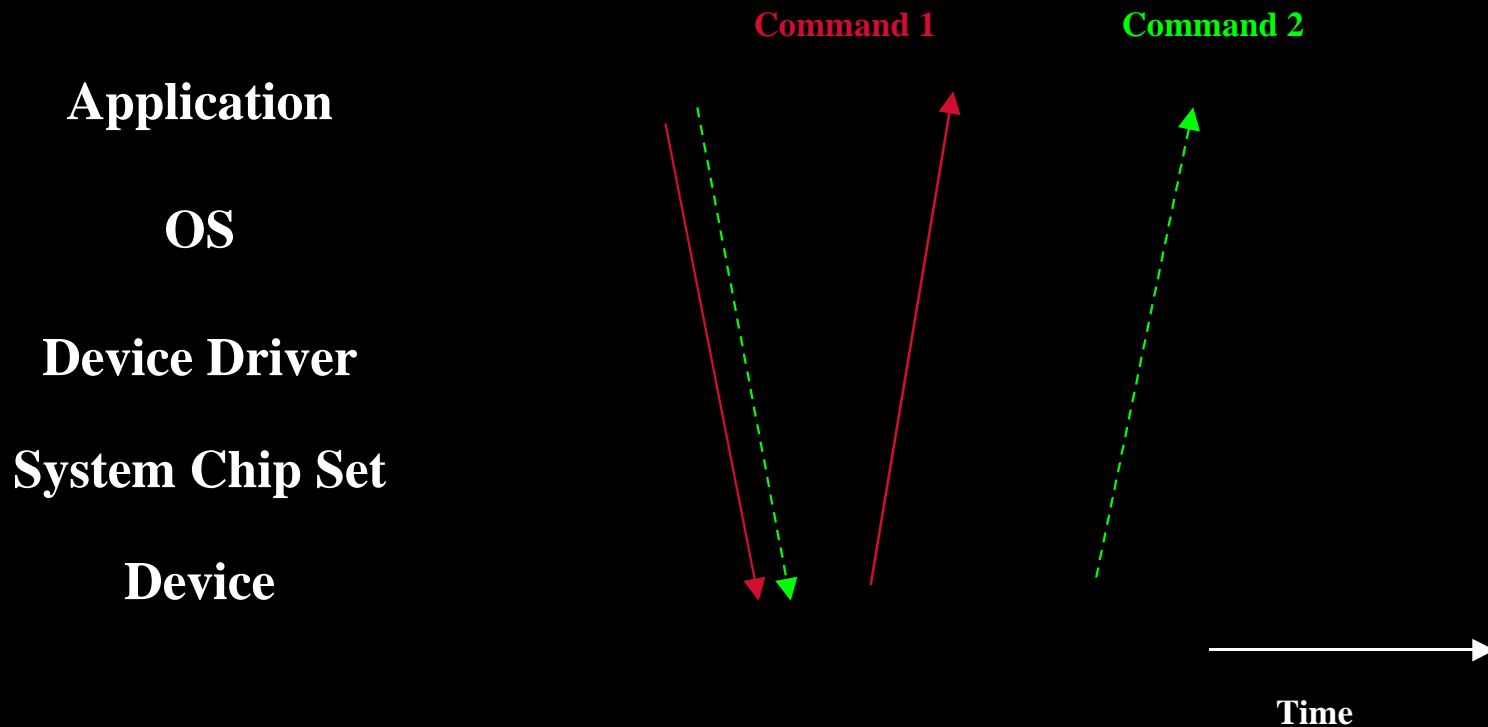
# Command Overhead

- Key is to avoid the long turn around time between the end of one command and the beginning of the next command.



# Command Overhead

- Queuing at the device level can eliminate the remaining host command overhead, and remove some complexity (and cost) from the system chip.





## Net Host Transfer Rate

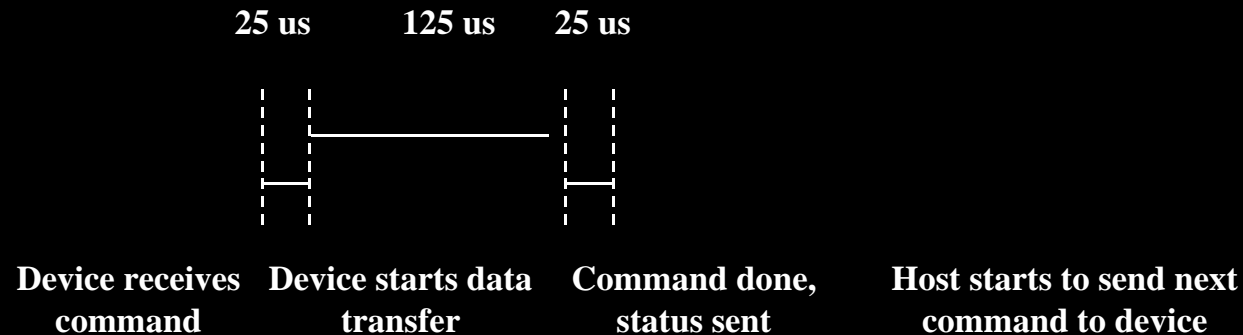
- Implementing both faster transfer rates and queuing to reduce host overhead improves net host transfer rate.

Previous example of 4K DMA read (prefetch hit)

32 MB/s & old overhead -> 10 MB/s

16 MB/s & new overhead -> 13 MB/s

32 MB/s & new overhead -> 22 MB/s



- Faster burst transfer rates are essential (net < burst)

high command overhead can be amortized over larger transfers (e.g. 32 MB/s and 250 us requires at least 11K to net 20 MB/s).

## **Net Host Transfer Rate**

---

- **Clearly at 16 MB/s burst transfer rate a 16 MB/s net disk transfer rate in 1997 cannot be sustained.**
- **At 1997 disk rates and 22 MB/s burst host transfer rate the host request would have to be minimum of 16 KBytes.**
- **At 1999 disk rates and 32 MB/s burst host transfer rate the host request would have to be a minimum of 32 KBytes.**
- **At 2001 disk rates and 64 MB/s burst host transfer rate the host request would have to be a minimum of 64 KBytes.**

## Command Overlapping

---

### • Interrupts

All overlap capable devices are assigned a device id (0-3) by the system during configuration.

When an interrupt is requested by a device for PIO data, DMA data, or status transfer, it toggles a single, shared INTRQ line.

When system chip detects the IDE bus is free it issues a read of a control register address responded to by all overlap capable devices (legacy devices naturally ignore it).

Device ID N toggles DB(N) to indicate its interest.

System selects device to start transfer, which in the case of queuing must be preceded by a read of the queue tag number.

### • Ending Interrupts

System chip set can clear the bus at any time (e.g. not do a PIO, deassert DMACK) to service other interrupts, in which case the process is repeated.

## Command Overlapping

---

- **Issuing commands**

**System chip set must allow host to issue another command regardless of what it is currently doing (e.g. transferring data on the IDE bus).**

**System chip set halts current IDE bus activity and uses combinations of CS0 and CS1, which are not currently legal, select the device in a quick and simple manner.**

**Device must be able to respond to reads/writes to control and command registers regardless of what it is currently doing.**

- **Legacy considerations**

**All devices boot with overlap (and queuing) disabled.**

**System responsible for enabling overlap/queuing.**

**Legacy devices can be accessed if and only if there is no command outstanding for an overlap capable device.**

## Cost Reduction

---

- System chip sets are IO limited, so costs are reduced by reducing pins.

Six pins are commonly duplicated for each IDE port: DMACK, DMREQ, INTRQ, IORDY, DIOR, DIOW.

4 (DMACK, DMREQ, INTRQ, IORDY) could easily be saved.

Duplicate DIOR and DIOW could be saved, although a new pin would be needed to select port 0 and port 1 devices separately.

Since DIOR and DIOW are high speed signals, keeping them duplicated and reducing the cable length they are propagated through might be best.

- Some more silicon is required (INT state machine), but that comes for free today.
- Number of legacy devices limited to 2 out of 4 if the pin savings are realized.