

CS152 – Computer Architecture and Engineering

Lecture 21 – Buses and Networks

2003-11-06

Dave Patterson
(www.cs.berkeley.edu/~patterson)

www-inst.eecs.berkeley.edu/~cs152/



CS 152 L21 Buses & Networks (1)

Patterson Fall 2003 © UCB

Review #1 / 2 Things to Remember

- Virtual memory to Physical Memory Translation too slow?
 - Add a cache of Virtual to Physical Address Translations, called a **TLB**
 - Need more compact representation to reduce memory size cost of simple 1-level page table (especially 32- → 64-bit address)
- Spatial Locality means Working Set of Pages is all that must be in memory for process to run fairly well
- Virtual Memory allows protected sharing of memory between processes with less swapping to disk



CS 152 L21 Buses & Networks (2)

Patterson Fall 2003 © UCB

Review #2 / 2: TLB/Virtual Memory

- VM allows many processes to share single memory without having to swap all processes to disk
- **Translation, Protection, and Sharing** are more important than memory hierarchy
- Page tables map virtual address to physical address
 - TLBs are a cache on translation and are extremely important for good performance
 - Special tricks necessary to keep TLB out of critical cache-access path
 - TLB misses are significant in processor performance:
 - These are funny times: most systems can't access all of 2nd level cache without TLB misses!



CS 152 L21 Buses & Networks (3)

Patterson Fall 2003 © UCB

Administrivia

- Design Doc for Final Project due
 - Thur 11/6 if finished lab 6 Friday or Monday
 - Mon 11/10 if finished lab 6 Tuesday or later
- Tue 11/11: Veteran's Day (no lecture)
- Fri 11/14: Demo Project modules
- Wed 11/19: 5:30 PM Midterm 2 in 1 LeConte
 - No lecture Thu 11/20 due to evening midterm
- Tues 11/22: Field trip to Xilinx
- CS 152 Project Week: 12/1 to 12/5
 - Mon: TA Project demo, Tue: 30 min Presentation, Wed: Processor races, Thu: lecture, Fri: Report




CS 152 L21 Buses & Networks (4)

Patterson Fall 2003 © UCB

What is a bus?

A Bus Is:

- shared communication link
- single set of wires used to connect multiple subsystems



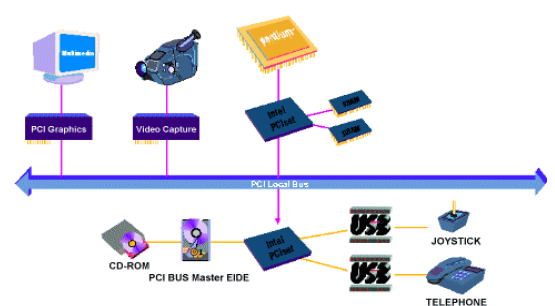
- A Bus is also a fundamental tool for composing large, complex systems
 - systematic means of abstraction



CS 152 L21 Buses & Networks (5)

Patterson Fall 2003 © UCB


Buses: PCI



CS 152 L21 Buses & Networks (6)

Patterson Fall 2003 © UCB

Advantages of Buses




- Versatility:
 - New devices can be added easily
 - Peripherals can be moved between computer systems that use the same bus standard
- Low Cost:
 - A single set of wires is shared in multiple ways



CS 162 L21 Buses & Networks (7)

Patterson Fall 2003 © UCB

Disadvantage of Buses




- It creates a communication bottleneck
 - The bandwidth of that bus can limit the maximum I/O throughput
- The maximum bus speed is largely limited by:
 - The **length** of the bus
 - The **number** of devices on the bus
 - The need to support a range of devices with:
 - Widely varying latencies
 - Widely varying data transfer rates



CS 162 L21 Buses & Networks (8)

Patterson Fall 2003 © UCB

The General Organization of a Bus



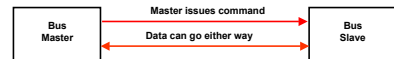
- Control lines:
 - Signal requests and acknowledgments
 - Indicate what type of information is on the data lines
- Data lines carry information between the source and the destination:
 - Data and Addresses
 - Complex commands



CS 162 L21 Buses & Networks (9)

Patterson Fall 2003 © UCB

Master versus Slave



- A **bus transaction** includes two parts:
 - Issuing the command (and address) – request
 - Transferring the data – action
- Master is the one who starts the bus transaction by:
 - issuing the command (and address)
- Slave is the one who responds to the address by:
 - Sending data to the master if the master ask for data
 - Receiving data from the master if the master wants to send data



CS 162 L21 Buses & Networks (10)

Patterson Fall 2003 © UCB

Types of Buses


- Processor-Memory Bus (design specific)
 - Short and high speed
 - Only need to match the memory system
 - Maximize memory-to-processor bandwidth
 - Connects directly to the processor
 - Optimized for cache block transfers
- I/O Bus (industry standard)
 - Usually is lengthy and slower
 - Need to match a wide range of I/O devices
 - Connects to the processor-memory bus or backplane bus
- Backplane Bus (standard or proprietary)
 - Backplane: an interconnection structure within the chassis
 - Allow processors, memory, and I/O devices to coexist
 - Cost advantage: one bus for all components



CS 162 L21 Buses & Networks (11)

Patterson Fall 2003 © UCB

A Computer System with 1 Bus: Backplane Bus




- A single bus (the backplane bus) is used for:
 - Processor to memory communication
 - Communication between I/O devices and memory
- Advantages: Simple and low cost
- Disadvantages: slow and the bus can become a major bottleneck
- Example: IBM PC - AT



CS 162 L21 Buses & Networks (12)

Patterson Fall 2003 © UCB

A Two-Bus System




- I/O buses tap into the processor-memory bus via bus adaptors:
 - Processor-memory bus: mainly for processor-memory traffic
 - I/O buses: provide expansion slots for I/O devices
- Apple Macintosh-II
 - NuBus: Processor, memory, and a few selected I/O devices
 - SCCI Bus: the rest of the I/O devices



CS 162 L21 Buses & Networks (13)

Patterson Fall 2003 © UCB

A Three-Bus System (+ backside cache)



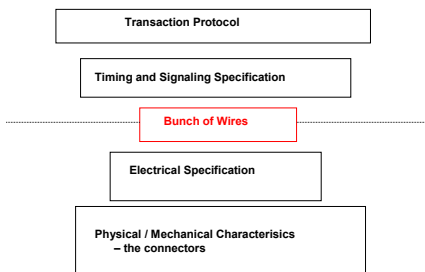
- A small number of backplane buses tap into the processor-memory bus
 - Processor-memory bus is only used for processor-memory traffic
 - I/O buses are connected to the backplane bus
- Advantage: loading on the processor bus is greatly reduced



CS 162 L21 Buses & Networks (14)

Patterson Fall 2003 © UCB

What defines a bus?

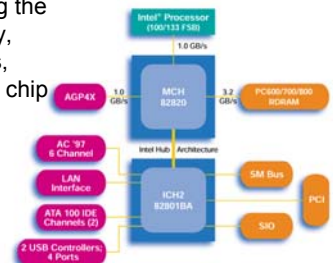


CS 162 L21 Buses & Networks (15)

Patterson Fall 2003 © UCB

Main components of Intel Chipset: Pentium III

- Northbridge: a DMA controller, connecting the processor to memory, the AGP graphic bus, and the south bridge chip
- Southbridge: I/O
 - PCI bus
 - Disk controllers
 - USB controllers
 - Audio
 - Serial I/O
 - Interrupt controller
 - Timers




CS 162 L21 Buses & Networks (16)

Patterson Fall 2003 © UCB

What is DMA (Direct Memory Access)?

- Typical I/O devices must transfer large amounts of data to memory of processor:
 - Disk must transfer complete block
 - Large packets from network
 - Regions of frame buffer
- DMA gives external device ability to access memory directly: much lower overhead than having processor request one word at a time.
- Issue: Cache coherence:
 - What if I/O devices write data that is currently in processor Cache?
 - The processor may never see new data!
 - Solutions:
 - Flush cache on every I/O operation (expensive)
 - Have hardware invalidate cache lines (remember "Coherence" cache misses?)

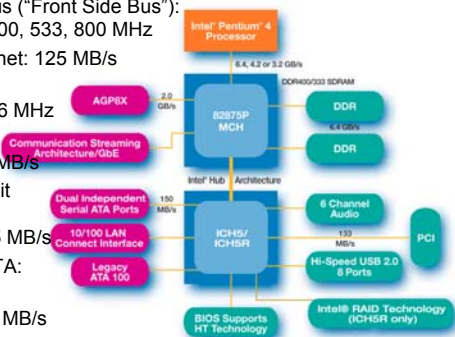


CS 162 L21 Buses & Networks (17)

Patterson Fall 2003 © UCB

Main components of Intel Chipset: Pentium 4

- System Bus ("Front Side Bus"): 64 bits x 400, 533, 800 MHz
- Gbit Ethernet: 125 MB/s
- Hub bus: 8 bits x 266 MHz
- 2 Serial ATA: 150 MB/s
- 10/100 Mbit Ethernet: 1.25 - 12.5 MB/s
- Parallel ATA: 100 MB/s
- 8 USB: 60 MB/s
- 1 PCI: 32b x 33 MHz



CS 162 L21 Buses & Networks (18)

Patterson Fall 2003 © UCB

I/O Chip Sets Customize Processor to App

Target Segment	875P Chip set	845GL Chip set
System Bus (64 bit)	Performance PC 800/533 MHz	Value PC 400 MHz
Memory Controller Hub ("North bridge")		
Package size, pins	42.5 x 42.5 mm, 1005	37.5 x 37.5 mm, 760
Memory Speed	DDR 400/333/266 SDRAM	DDR 266/200, PC133 SDRAM
Memory buses, widths	2 x 72	1 x 64
Maximum Memory Capacity	4 GB	2 GB
Memory Error Correction available?	Yes	No
AGP Graphics Bus, Speed	Yes, 8X or 4X	No
Graphics controller	External	Internal (Extreme Graphics)
CSA Gigabit Ethernet interface	Yes	No
South bridge interface speed (8 bit)	266 MHz	266 MHz
I/O Controller Hub ("South bridge")		
Package size, pins	31 x 31 mm, 460	31 x 31 mm, 421
PCI bus: width, speed, masters	32-bit, 33 MHz, 6 masters	32-bit, 33 MHz, 6 masters
Ethernet MAC controller, interface	100/10 Mbit	100/10 Mbit
USB 2.0 ports, controllers	8, 4	6, 3
ATA 100 ports	2	2
Serial ATA 150 controller, ports	Yes, 2	No
RAID 0 controller	Yes	No
AC-97 audio controller, interface	Yes	Yes
Management	SMbus 2.0, GPIO	SMbus 2.0, GPIO

CS 162 L21 Buses & Networks (19)

Patterson Fall 2003 © UCB

Networks

Networks are major medium used to communicate between computers. Key characteristics of typical networks:

- Distance:** 0.01 to 10,000 kilometers
Local Area Network (LAN) <1 km vs.
Wide Area Network (WAN) to 10000 km
- Speed:** 0.001 MB/sec to 100 MB/sec
- Topology:** Bus, ring, star, tree
- Shared lines:** None (switched point-to-point) or shared (multidrop)

CS 162 L21 Buses & Networks (20)

Patterson Fall 2003 © UCB

Protocols: HW/SW Interface

- Internetworking:** allows computers on independent and incompatible networks to communicate reliably and efficiently;
 - Enabling technologies: SW standards that allow reliable communications without reliable networks
 - Hierarchy of SW layers, giving each layer responsibility for portion of overall communications task, called **protocol families** or **protocol suites**
- Transmission Control Protocol/Internet Protocol (TCP/IP)**
 - This protocol family is the basis of the Internet
 - IP makes best effort to deliver; TCP guarantees delivery
 - TCP/IP used even when communicating locally: NFS uses IP even though communicating across homogeneous LAN

CS 162 L21 Buses & Networks (21)

Patterson Fall 2003 © UCB

Protocol

- Key to **protocol families** is that communication occurs **logically** at the same level of the protocol, called **peer-to-peer**, but is **implemented via services at the lower level**
- Danger is each level increases latency if implemented as hierarchy (e.g., multiple check sums)

CS 162 L21 Buses & Networks (22)

Patterson Fall 2003 © UCB

Open Systems Interconnect (OSI)

- Open Systems Interconnect (OSI) developed a model that popularized describing networks as a series of 7 layers


The OSI Layer Model	
OSI	TCP/IP
Layer 7 Application	Application Telnet, FTP, NFS, NFS
Layer 6 Presentation	Session e.g. RPC
Layer 5 Session	Transport Socket/Sockets - TL1
Layer 4 Transport	TCP UDP
Layer 3 Network	Network IP - ARP/RARP/ICMP
Layer 2 Data Link	Physical Protocol Ethernet/TSP/HD/PPP
Layer 1 Physical	Transmission medium Coax, Fiber, 10baseT

CS 162 L21 Buses & Networks (23)

Patterson Fall 2003 © UCB

TCP/IP packet


- Application sends message
- TCP breaks into 64KB segments, adds 20B header
- IP adds 20B header, sends to network
- If Ethernet, broken into 1500B packets with headers, trailers
- Header, trailers have length field, destination, window number, version, ...



CS 162 L21 Buses & Networks (24)

Patterson Fall 2003 © UCB

FTP From Stanford to Berkeley



- BARRNet is WAN for Bay Area
- T1 is 1.5 mbps leased line; T3 is 45 mbps; FDDI is 100 mbps LAN
- IP sets up connection, TCP sends file

Long Haul Networks (or WANs)

- 10 km to 10,000 km
- **packet-switch**: At each hop, a packet is stored (for recovery in case of failure) and then forwarded to the proper target according to the address in the packet.
- Destination systems reassembles packets into a message.
- Most networks today use packet switching, where packets are individually routed from source to destination.

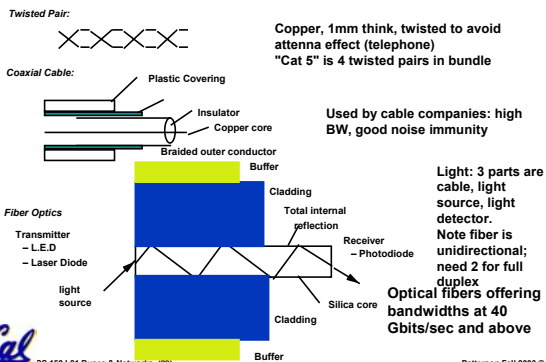
Connecting Networks

- **Routers or Gateways**: these devices connect LANs to WANs or WANs to WANs and resolve incompatible addressing.
 - Generally slower than bridges, they operate at the **internetworking protocol (IP) level**: OSI layer 3
 - Routers divide the interconnect into separate smaller subnets, which simplifies manageability and improves security
- **Bridges**: connect LANs together, passing traffic from one side to another depending on the addresses in the packet
 - operate at the **Ethernet protocol level**: OSI layer 2
 - usually simpler and cheaper than routers
- **Hubs**: extend multiple segments into 1 LAN.
 - Only transmit one message can at a time
 - operate at the **Physical level**: OSI layer 1

Local Area Networks: Ethernet

- Ethernet packets vary 64 to 1518 Bytes
- Ethernet link speed available at 10M, 100M, and 1000M bits/sec, with 10,000M bits/sec available soon
- Although 10M and 100M bits/sec can share the media with multiple devices, 1000M bits/sec and above relies on point-to-point links and switches

Network Media



Wireless Local Area Networks

- IEEE 802.11("WiFi") extended Ethernet to communicate through the air. 3 variations:
 - 802.11b, peak of 11 Mbits/second
 - 802.11a, peak of 54 Mbits/second
 - 802.11g, peak of 22 Mbits/second
- In practice, the delivered rates in the field are about a third of the peak rates in the lab.
- It replaces the bottom layers of the OSI standard, which Ethernet labels the MAC layer and PHY layer, with radio

Radio Overview

- A radio wave is an electromagnetic wave propagated by an antenna
- Radio waves are modulated: sound signal is superimposed on stronger radio wave that carries the data ("carrier signal")
- 802.11b and 802.11g use 2.4 GHz carrier and 802.11a uses 5 GHz frequency carrier.
 - Both actually use small % of frequencies on either side of the norm => giving them multiple channels on which to transmit.
 - If two transmitters collide, they hop to another channel and try again



CS 162 L21 Buses & Networks (31)

Patterson Fall 2003 © UCB

Radio Overview

- Bit error rate (BER) of wireless link is determined by received signal power, noise due to interference caused by the receiver hardware and interference from other sources
 - Noise typically proportional to radio frequency BW



CS 162 L21 Buses & Networks (32)

Patterson Fall 2003 © UCB

Wireless Network Challenges

- 1) Devices are mobile or wiring is inconvenient, which means the wireless network must rearrange itself dynamically
- 2) Wireless signals are not protected => subject to mutual interference, especially as devices move, and to eavesdropping
- 3) Power: both because mobile devices tend to be battery powered and because antennas radiate power to communicate and little of it reaches the receiver
 - Raw bit error rates typically 1,000 to 1,000,000 times higher than copper wire



CS 162 L21 Buses & Networks (33)

Patterson Fall 2003 © UCB

2 primary architectures for wireless networks

- **Base stations** connected by wire for longer-distance communication, and mobile units communicate only with a single local base station (802.11)
- **Peer-to-peer** architectures allow mobile units to communicate with each other, and messages hop from one unit to the next until delivered to the desired unit
- peer-to-peer more reconfigurable, but base stations more reliable since only 1 hop between the device and the station



CS 162 L21 Buses & Networks (34)

Patterson Fall 2003 © UCB

Peer Instruction

- Ethernet packet size is 64 to 1538 Bytes
- If you could redesign packets just for wireless, how would they look
 - 1) Due to the higher Bit Error Rate of wireless, you would like smaller packets
 - 2) Ethernet was inspired by Aloha net which was a wireless network, so that packet sizes are fine as is
 - 3) To get greater bandwidth when using air as the medium, you'd like larger packets



CS 162 L21 Buses & Networks (35)

Patterson Fall 2003 © UCB

Smaller packets yet Ethernet?

- 802.11 allows MAC layer to fragment large messages into several smaller messages
- The MAC layer of the receiving device then reassembles these smaller messages into the original full Ethernet message



CS 162 L21 Buses & Networks (36)

Patterson Fall 2003 © UCB

Privacy yet Radio?

- 802.11 offers “Wired Equivalent Privacy”
- It uses a pseudo-random number generator initialized by a shared secret key.
- Operators initialize access points and end-user stations with the secret key.
- A pseudo-random sequence of bits equal to the largest packet is combined with the real packet to encode the packet transmitted in air.



802.11 vs. Cellular Telephony

- Which is cheaper?
- Why?
 - Distance?
 - Universal access?
 - Voice vs. Data?
 - Automobile?
 - Internet vs. Telephone infrastructure?



Peer Instruction

- Which of the following are true?
- 1) Protocol stacks are an example of using abstraction to hide complexity.
 - 2) TCP/IP is used for WANs, but LANs use a protocol stack appropriate for the lower latency and higher bandwidths.
 - 3) Although the 802.11 LAN standard is wireless like the cell phone, there is little commonality between the two technologies.



Summary

- Buses are an important technique for building large-scale systems
 - Their speed is critically dependent on factors such as length, number of devices, etc.
 - Critically limited by capacitance
- Direct Memory Access (dma) allows fast, burst transfer into processor's memory:
 - Processor's memory acts like a slave
 - Probably requires some form of cache-coherence so that DMA'ed memory can be invalidated from cache.
- Networks and switches popular for LAN, WAN
- Networks and switches starting to replace buses on desktop, even inside chips

