

Designing a bus system

- A bus need to support
 - cache-memory traffic
 - I/O-memory traffic
 - Processor-I/O traffic
- The first one depends on cache miss rate and replacement
- The number of cycles for each transaction is to read a new line or write a dirty line back
- For disk, each disk controller may support many disks
- Disk controller is busy to initiate a transfer and to transfer data to/from memory for the actual data transfer (as opposed to whole transaction) operation (so 1-2 ms out of 15 ms or so in our earlier example)
- Bus is only busy during actual transfer
- Disk controller may transfer in burst mode (multiple bytes in one transaction)

1

Designing a bus system (Continued)

- The bus should not be designed to keep it busy 100% of the time
- Suppose a bus takes 200 processor cycles to transfer a 4-word (16-bytes) block in and out of memory (assume a 200 MHz processor)
- Then it will take $(4K/16)*200 = 50,000$ cycles (it is really is a bit more, but we are simplifying) to transfer a 4K byte block
- Processor and controller may take an additional 50,000 cycles to establish a transfer and complete it
- They may use the bus for 10,000 cycle
- A 4KB transfer keeps disk busy for 15ms (3,000,000 cycles)
- A 4KB transfer keeps disk controller busy for 100,000 cycles
- A 4KB transfer keeps processor busy for 50,000 cycles
- A 4KB transfer keeps bus busy for 60,000 cycles
- A New line (32 bytes) fetch keeps the bus busy for 200 cycles
- A dirty line (32 bytes) write keeps the bus busy for 400 cycles

2

Designing a bus system (Continued)

- Each disk can support $1\text{sec}/15\text{ms} = 66$ 4KB transfers/sec
- Each controller can support $200M/100,000 = 2000$ transfers/sec
- Each controller can support $2000/66 = 30$ disks
- A processor can support $200M/50,000 = 4000$ transfers/sec
- However, the processor should not be busy with disk only
- With 25% processor capacity, it can only support 1000 transfers/sec
- Or number of disk that can be kept busy is $1000/66 = 15$
- The bus can support $200M/60,000 = 3333$ transfers/sec
- Or it can support $3333/66 = 50$ disks
- However, the bus should not be loaded, say, more than 25% times with disk load, so it can really support only 12 disks
- The number of disks is decided based on the critical resource
- BUS HAS TO SUPPORT CACHE TRAFFIC BASED ON MISS RATE

3

Exceptions

- Exceptions are just that – Changes in the normal execution of a program
- Two types of exceptions
 - External Condition: I/O interrupt, power failure, user termination signal (Ctrl-C)
 - Internal Condition: Bad memory read address (not a multiple of 4), illegal instructions, overflow/underflow.
- Interrupts – external
- Exceptions – internal
- Usually we can refer to both by the general term "Exception" though.
- In either case, we need some mechanism by which we can handle the exception generated.

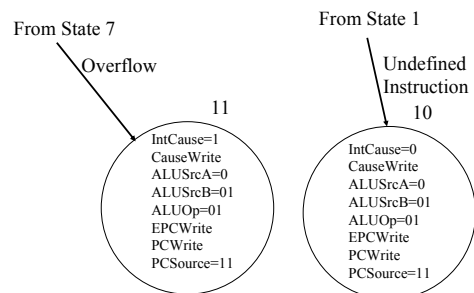
4

How Exceptions are Handled

- We need two special registers
 - EPC: 32 bit register to hold address of current instruction
 - Cause: 32 bit register to hold information about the type of exception that has occurred.
- Simple Exception Types
 - Undefined Instruction
 - Arithmetic Overflow
- Another type is Vectored Interrupts
 - Do not need cause register
 - Appropriate exception handler jumped to from a vector table

5

Two new states for the Multi-cycle CPU



6

Exceptions and Pipelined Processors

- The pipelined processor must also have some mechanism to handle exceptions
- Fortunately, the pipelining will not make exceptions much more difficult
- Look at p. 506 for datapath with exception handling
 - Added 40000040 as another input for PC (the exception handler address)
 - EPC and Cause register are added to datapath

7

Virtual Memory and Exceptions

- Virtual Memory TLB Misses
 - Page is just not in TLB
 - Bring page information into TLB
 - Page is not in Main Memory
 - Page Fault requires OS to intervene
 - Exception – Page Fault

8

Handling a Page Fault

- 1) Look up the page table entry corresponding to the virtual address to find the location of the referenced page on disk
 - 2) Choose a page in main memory to replace
 - If that page has been written to in the past (dirty bit is set)
 - Recopy the page back to the disk
 - 3) Move the new page into main memory from the disk
-
- Step 2 may be very slow if page to be replaced is dirty
 - Step 3 will take millions of clock cycles to complete
 - So push this process to the side temporarily and do other meaningful work
 - Then later we can return from the exception handler and continue the program execution

9

Exceptions in the Exception Handler

- Problem: What if another exception occurs within the exception handler itself?
 - Impossible to return to initial exception location, since EPC will be overwritten
- Solution: Have the ability to turn off exception handling.
 - Set a bit that can disable other exceptions from affecting execution

10

I/O Devices and Exceptions

- I/O devices will generate interrupts to notify the processor
- Who will handle these interrupts?
 - Operating System
 - Provides interface to system I/O devices, so you don't need to do all low-level operations
 - Provide some fairness in resource usage, as well as scheduling to improve throughput
- Memory Mapped I/O versus Dedicated I/O instructions

11

Communication with I/O Devices

- Reading and Writing to devices usually requires several steps
- Status Registers
 - Hold information pertaining to the state of the device
 - Done bit or Error bit, etc.
 - May also be written to for notifying device when the data input is ready
- Data Registers
 - Buffers for information
 - Examples: character to be printed, Ethernet packet, etc.
 - Some are only readable, others are only writeable. Sometimes they are both R/W.
- User I/O is managed by the supervisor (kernel) level, since the device address space is not usually available to a user

12

Polling versus Interrupt-driven I/O

- Polling
 - Processor must check whether or not I/O device has new meaningful information
 - Large overhead costs
 - Still sees some use though with very slow devices that are routinely used (e.g. mouse)
- Interrupt-driven I/O
 - I/O device will notify processor by way of interrupt to request services
 - Not synchronous with instructions
 - Vectored interrupts or EPC
 - Can have various interrupt levels to show priority

13

Direct Memory Access (DMA)

- Memory/Device data transfers without constant use of the processor
- DMA is the bus master, thus it directs the traffic
- DMA Transfer
 - Processor must inform DMA of operation to perform along with the various parameters (e.g. device address, source address, destination address, bytes to transfer, ...)
 - DMA starts the transfer and controls the bus, performing the requested operation
 - When the operation is done, the DMA controller sends an interrupt to the CPU to let it know the status of the transfer
- There can be many DMA's in the same system
- Difficulties with virtual address translation
 - Coherency problem

14

MIPS Exception Related Information

- Coprocessor 0 is used for exceptions in MIPS
- P. A-32 in the textbook
- 4 registers accessible by *lwc0*, *mfc0*, *mtc0*, *swc0*
- MIPS uses \$k0 and \$k1 as kernel registers for exception handling

Register Name	Register Number	Usage
BadVAddr	8	Memory address at which the reference occurred
Status	12	Interrupt mask and enable bits
Cause	13	Exception type and pending interrupt bits
EPC	14	Address of exception-causing instruction

15

Review Material

- On Final Exam Key Points:
 - Number of cycles = $n + k - 1$ + bubbles
 - Forwarding unit does not insert bubbles
 - Hazard detection unit will insert bubble for anything that cannot be taken care of by forwarding
 - Design of data placement algorithms for efficient caching
 - Distinction between design issues vs. programming issues
 - Multi-way set associativity is design
 - Placing data appropriately is programming
- Memory Design will be a key issue
- Include caching, virtual memory
- And finally, I/O will be the major issue
- It is comprehensive
- Good Luck

16

Concluding Remarks

- Evolution vs. Revolution

“More often the expense of innovation comes from being too disruptive to computer users”

“Acceptance of hardware ideas requires acceptance by software people; therefore hardware people should learn about software. And if software people want good machines, they must learn more about hardware to be able to communicate with and thereby influence hardware engineers.”

17