OS Interfaces for I/O Systems — Today

- How does I/O hardware influence the OS?
- What I/O services does the OS provide?
- How does the OS implement those services?
- How can the OS improve the performance of I/O?

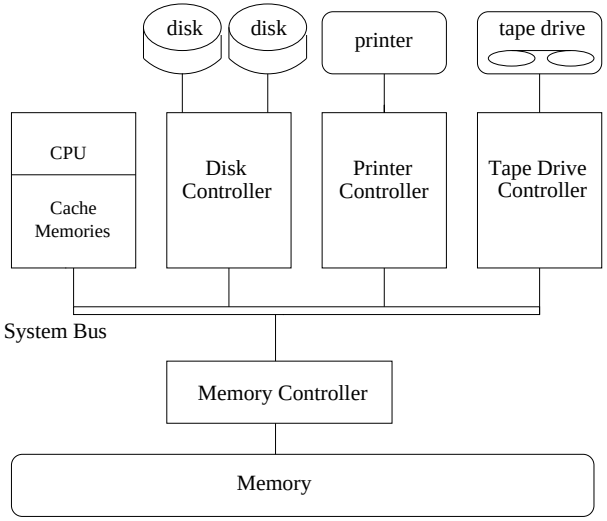

Architecture of I/O Systems

The hardware associated with an I/O device consists of four pieces:

- A **bus** that allows the device to communicate with the CPU, typically shared by multiple devices.
- A device **port** typically consisting of 4 registers:

  Status - whether the device is busy, whether data is ready, whether an error occurred.

  Control - command to perform

  Data-in - data being sent from the device to the CPU

  Data-out - data being sent from the CPU to the device
- A **controller** that can receive commands from the system bus, translate commands into device actions, and read data from and write data to the system bus.
- The device itself

Traditional devices: disk drive, printer, keyboard, modem, mouse, display

Non-traditional devices: joystick, robot actuators, flying surfaces of an airplane, fuel injection system of a car, ...

I/O Services Provided by OS

- Naming of files and devices. (On Unix, devices appear as files at the highest level of abstraction. The /dev directory contains file names, most of which actually represent devices.)
- Access control.
- Operations appropriate to the files and devices.
- Device allocation.
- Buffering, caching, and spooling to allow efficient communication with devices.
- I/O scheduling.
- Error handling and failure recovery associated with devices (command retries, for example).
- Device drivers to implement device-specific behaviors.

Communication between OS and I/O Device - Polling

Polling:

- CPU busy-waits until the status is idle.
- CPU sets the command register and data-out if it is an output operation.
- CPU sets the status to command-ready.
- Controller reacts to command-ready and sets the status to busy.
- Controller reads the command register and performs the command, placing a value in data-in if it is an input command.
- Assuming the operation succeeded, the controller changes the status to idle.
- CPU observes the change to idle and reads the data if it was an input operation.

- Good choice if data must be handled promptly, like for a modem or keyboard where data will be lost if it isn't moved from the device fast enough.
- What happens if the device is slow compared to the CPU?

Communication between OS and I/O Device - Interrupts

$\Longrightarrow$ Rather than have the CPU busy wait for a device to become available, the device can interrupt the CPU when it completes an I/O operation.

On an I/O interrupt:

- Determine which device caused the interrupt.
- If the last command was an input operation, retrieve the data from the device register.
- Start the next operation for that device.

Direct Memory Access

For devices that transfer large volumes of data at a time (like a disk block), it is expensive to have the CPU retrieve these one byte at a time.

Solution: Direct memory access (DMA)

- The device has a more sophisticated controller, a DMA controller, that can write directly to memory. Instead of data-in/data-out registers, it has an address register.
- The CPU tells the DMA the locations of the source and destination of the transfer.
- The DMA controller operates the bus and interrupts the CPU when the entire transfer is complete, instead of when each byte is ready.
- The DMA controller and the CPU compete for the memory bus, slowing down the CPU somewhat, but still providing better performance than if the CPU had to do the transfer itself.

Application Programmer's View of I/O Devices

The OS provides a high-level interface to devices, greatly simplifying the programmer's job.

- Standard interfaces are provided for related devices.
- Device dependencies are encapsulated in device drivers.
- New devices can be supported by an OS simply by providing a device driver with the device.

Device characteristics:

- Transfer unit: character or block
- Access method: sequential or random access
- Timing: synchronous or asynchronous. (Note that most devices are asynchronous, while I/O system calls are synchronous. The OS implements *blocking I/O*.)
- Sharable or dedicated
- Speed
- Operations: Input, output, or both

Examples:

- Keyboard?
- Disk?

I/O Buffering

I/O devices typically contain a small on-board memory where they can store data temporarily before transferring to/from the CPU.

- A disk buffer stores a block when it is read from the disk.
- It is transferred over the bus by the DMA controller into a buffer in physical memory.
- The DMA controller interrupts the CPU when the transfer is done.

Why buffer on the OS side?

- To cope with speed mismatches between device and CPU. Example: Compute the contents of a display in a buffer (slow) and then zap the buffer to the screen (fast).
- To cope with devices that have different data transfer sizes. Example: ftp brings the file over the network one packet at a time. Stores to disk happen one block at a time.
- To minimize the time a user process is blocked on a write. To write to a file, the write is done immediately to a kernel buffer, and control is returned to the user program. The write from the kernel buffer to the disk is done later.

Caching

Improve disk performance by reducing the number of disk accesses.

Idea: keep recently used disk blocks in main memory after the I/O call that brought them into memory completes.

Example: Read (diskAddress)
If (block in memory) return value from memory
Else ReadSector(diskAddress)

Example: Write (diskAddress)
If (block in memory) update value in memory
Else Allocate space in memory, read block from disk, and update value in memory

What should happen when we write to a cache?

- write-through policy (write to all levels of memory containing the block, including to disk). High reliability.
- write-back policy (write only to the fastest memory containing the block, write to slower memories and disk sometime later). Faster.

Putting the Pieces Together - a Typical Read Call

1. User process requests a read from a device.
2. OS checks if data is in a buffer. If not,
   (a) OS tells the device driver to perform input.
   (b) Device driver tells the DMA controller what to do and blocks itself.
   (c) DMA controller transfers the data to the kernel buffer when it has all been retrieved from the device.
   (d) DMA controller interrupts the CPU when the transfer is complete.
3. OS transfers the data to the user process and places the process in the ready queue.
4. When the process gets the CPU, it begins execution following the system call.

Summary

I/O is expensive for several reasons:

- I/O typically involves physical movement which is slow or network/phone line communication which is also slow.
- I/O devices are often contended for by multiple processes.
- I/O is typically supported via system calls and interrupt handling, which are slow.

Approaches to improving performance:

- Reduce the number of times data is copied by caching values in memory.
- Reduce the frequency of interrupts by using large data transfers when possible.
- Offload computation from the main CPU by using DMA controllers.
- Increase the number of devices to reduce contention for a single device and thereby improve CPU utilization.
- Increase physical memory to reduce amount of time paging and thereby improve CPU utilization.