ECS 150
Class Notes Jan. 8, 2004

Open Shop Model

Early computers, e.g., IBM 1620
Single user, hands-on, no operator, one job at a time, no operating system

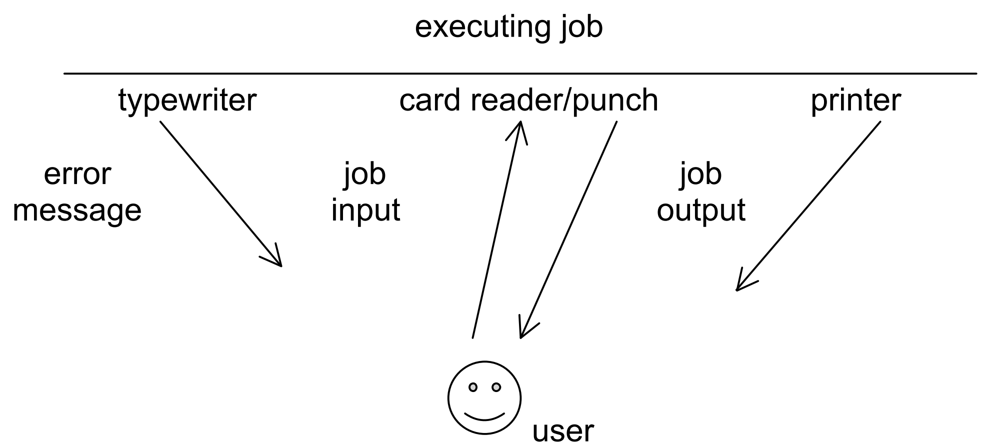

Utilization: Fraction of time used for computation

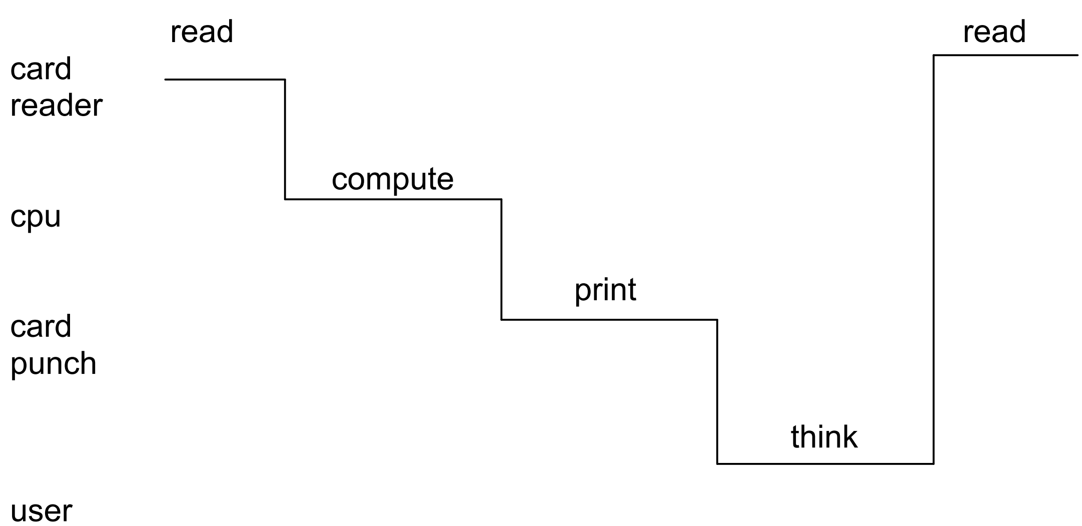

Operator-Driven Shop

Avoid cpu idle time
Operator loads jobs

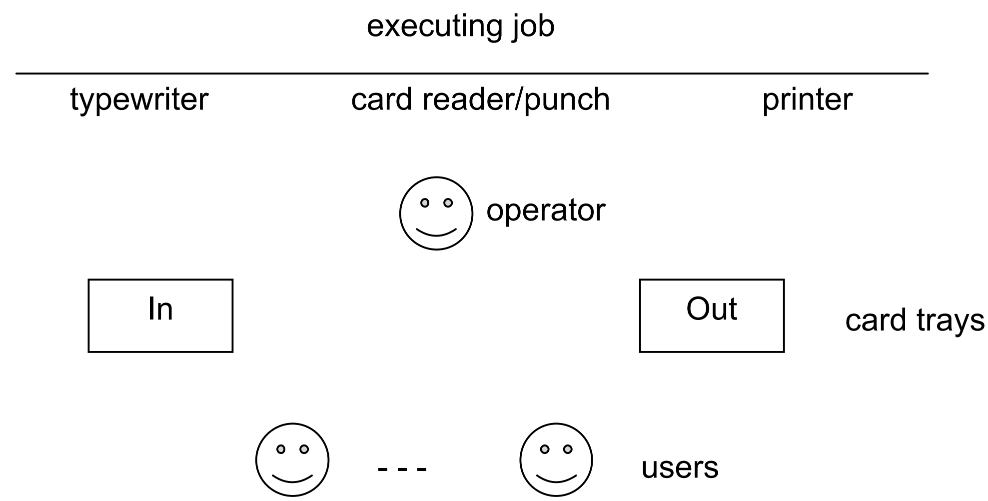

Users charged just for cpu time they use
<u>Batching</u> of similar job steps
<u>Priority</u> – users pay more for quick turnaround

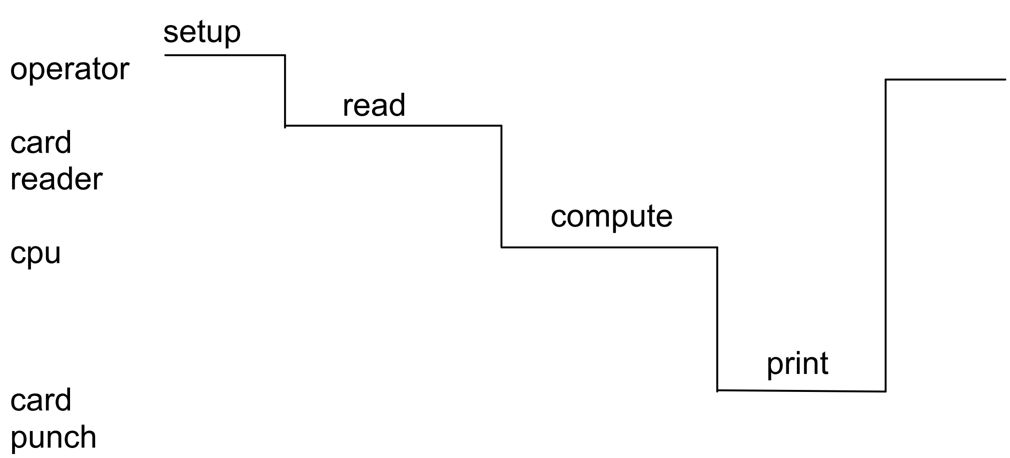

Offline Transport (I/O)

Automate I/O on separate (offline) computer, a.k.a., channel, satellite computer, peripheral processing unit (pppu)

Beginning of a real OS, called resident monitor
Reset machine after each job
load next job
accounting

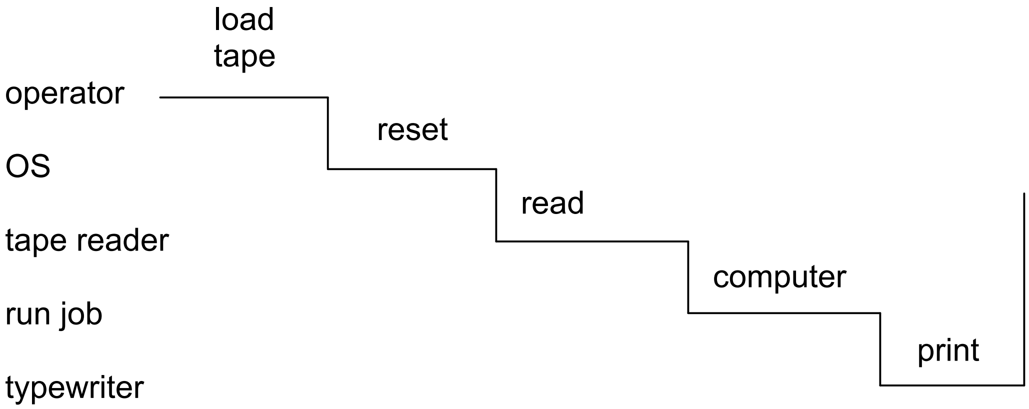

Spooling

I/O concurrently with running jobs
I/O devices generate interrups when I/O requests finished
Include disks in I/O device set, hold jobs
Add scheduler to OS

Utilization for spooling
Sufficient number of card readers and printers so always jobs ready to run
read/write to disks is faster
Computer spends a certain percent of time servicing interrupts for transport; not counted as useful computation time
OS resets machine between jobs

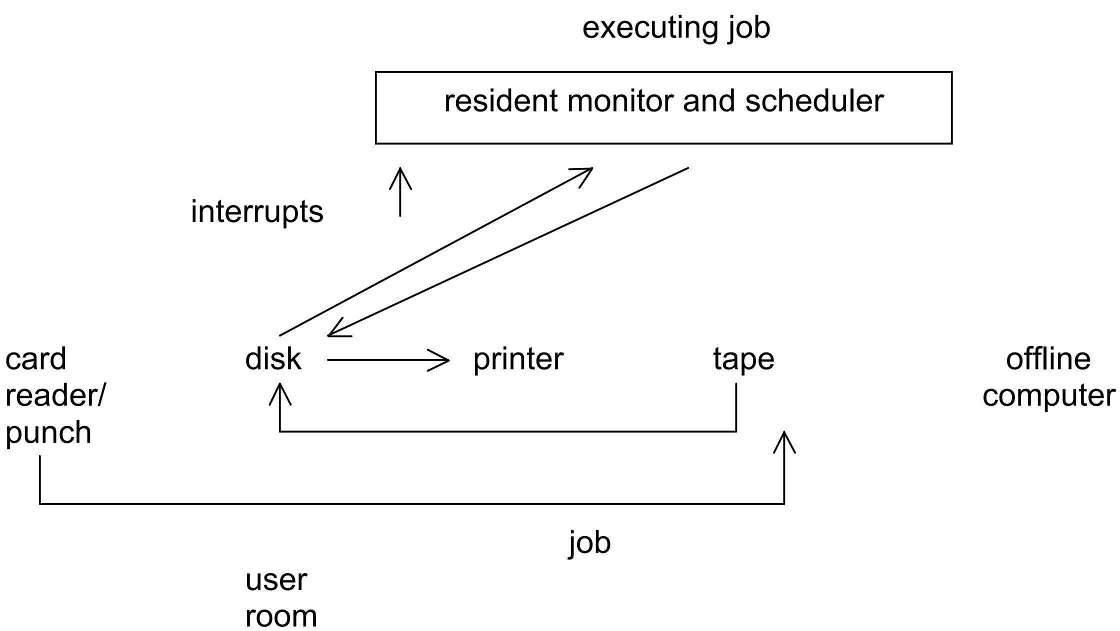

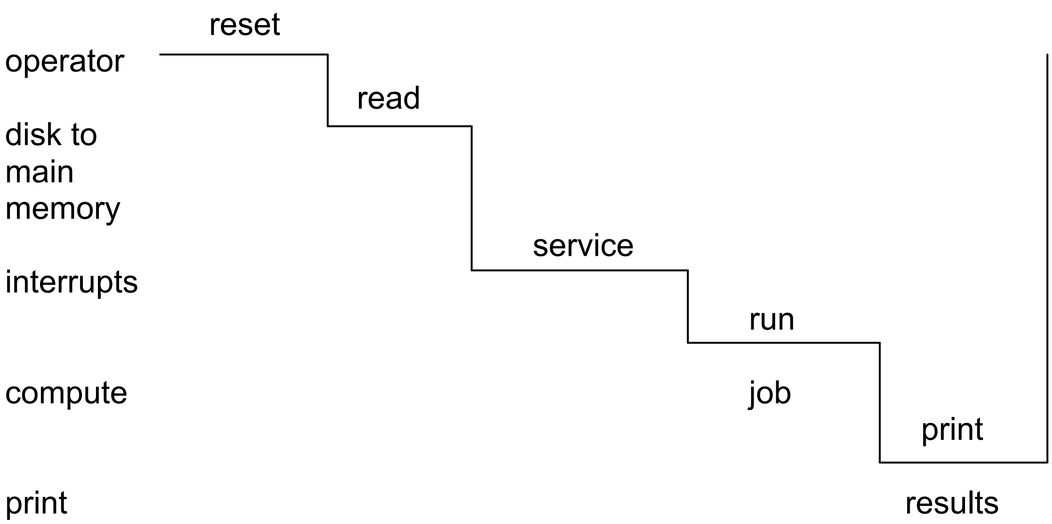

Note: Interrupts can occur at any time. I lumped the interrupt service activity into one interval.

Batch Multiprocessing/Multiprogramming

More than one job in main memory ready to run when current job performs an I/O request
Now a real OS that offers system calls

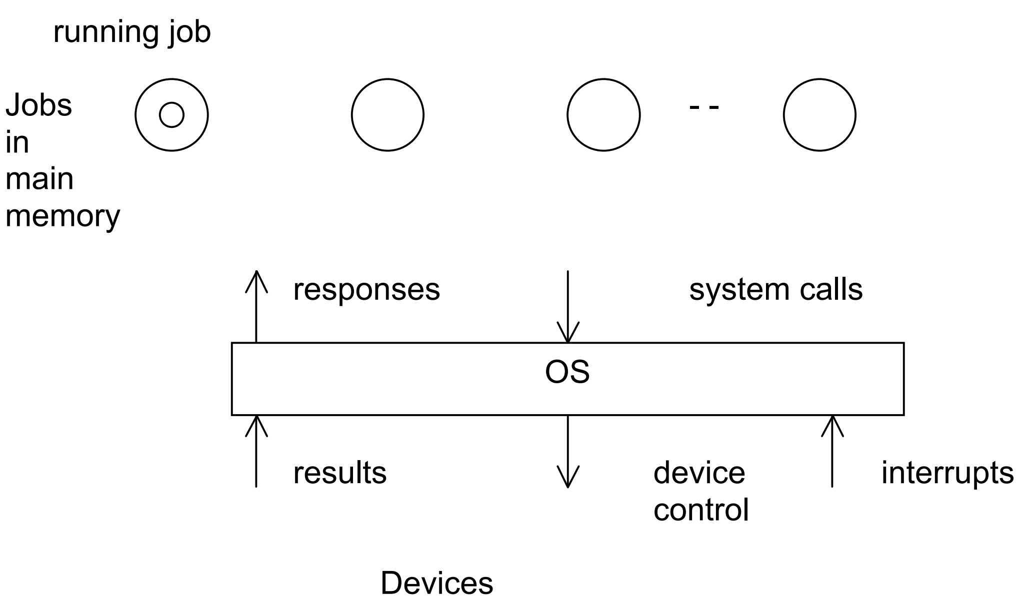

Life of a process:
Compute; I/O request; compute; I/O request; . . .

With only one job/process ready to run, CPU is often idle when I/O requests are serviced

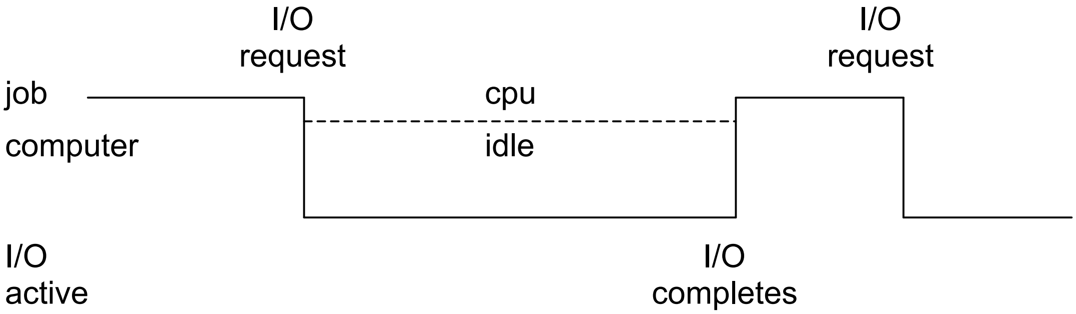

CPU is better utilized when computation can proceed concurrently with I/O

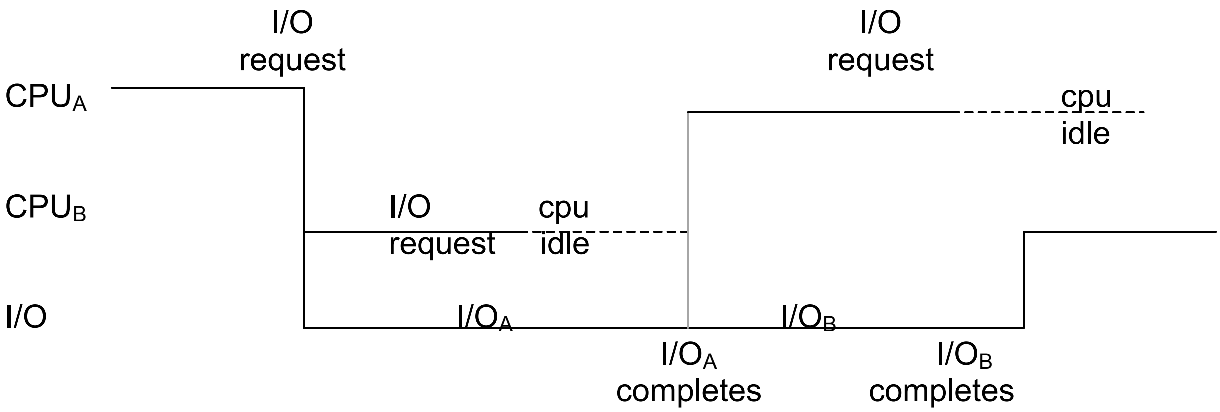

CPU idle time is reduced with more jobs in main memory; but how many jobs can main memory hold?