Disciplina
Sistemas de Computação

Aula 10

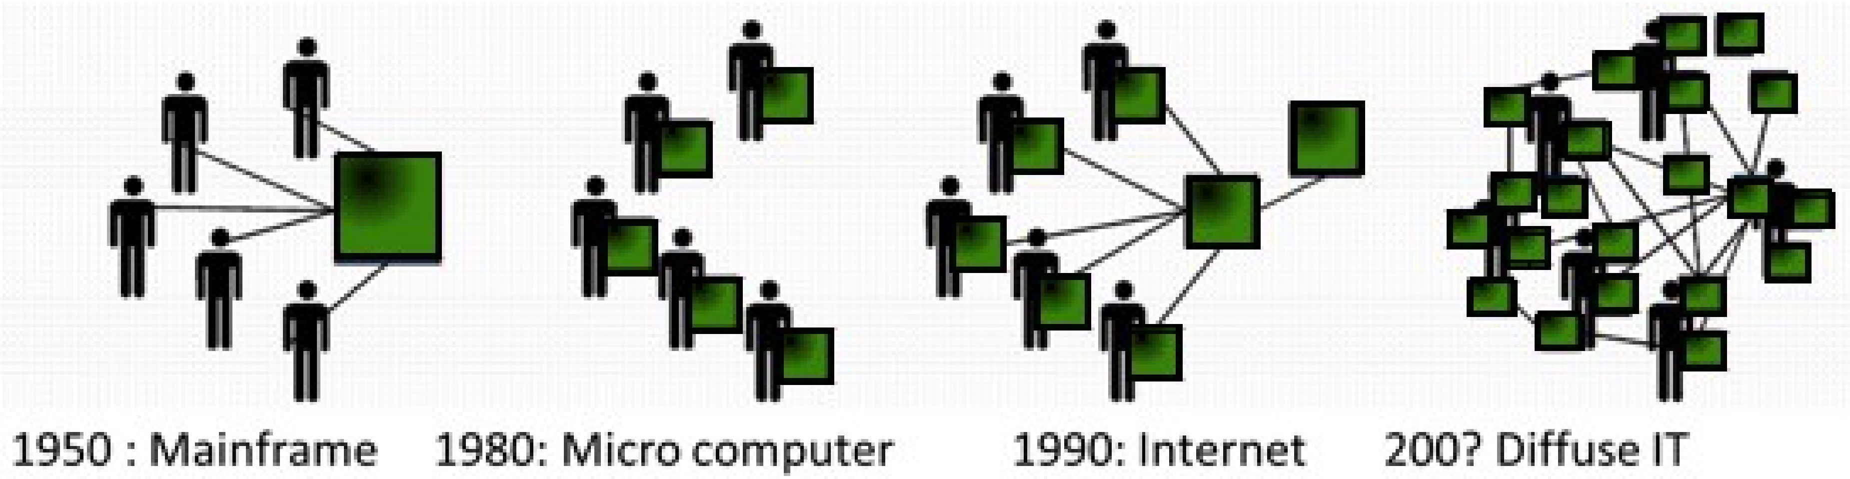

Short Term Scheduling

- The kernel runs the scheduler at least when
  - a process switches from running to waiting,
  - an interrupt occurs, or
  - a process is created or terminated.

- Non-preemptive system: the scheduler must wait for one of these events
- Preemptive system: the scheduler can interrupt a running process

Criteria for Comparing Scheduling Algorithms

- **CPU Utilization:** percentage of time that the CPU is busy.
- **Throughput:** number of processes completing in a unit of time.
- **Turnaround time:** length of time it takes to run a process from initialization to termination, including all the waiting time.
- **Waiting time:** total amount of time that a process is in the ready queue.
- **Response time:** time between when a process is ready to run and its next I/O request.

Scheduling Policies

- Ideally, choose a CPU scheduler that optimizes all criteria simultaneously (utilization, throughput,..), but this is not generally possible
- Instead, choose a scheduling algorithm based on its ability to satisfy a policy

Scheduling Policies

- Minimize average response time: provide output to the user as quickly as possible and process their input as soon as it is received.
- Minimize variance of response time: in interactive systems, predictability may be more important than a low average with a high variance.
- Maximize throughput: two components
  - minimize overhead (OS overhead, context switching)
  - efficient use of system resources (CPU, I/O devices)
- Minimize waiting time: give each process the same amount of time on the processor. This might actually increase average response time.

Scheduling Algorithms

- FCFS: First Come, First Served
- Round Robin: Use a time slice and preemption to alternate jobs.
- SJF: Shortest Job First
- Multilevel Feedback Queues: Round robin on each priority queue.
- Lottery Scheduling: Jobs get tickets and scheduler randomly picks winning ticket.

FCFS: First-Come-First-Served (or FIFO: First-In-First-Out)

- The scheduler executes jobs to completion in arrival order.
- In early FCFS schedulers, the job did not relinquish the CPU even when it was doing I/O.
- We will assume a FCFS scheduler that runs when processes are blocked on I/O, but that is non-preemptive, i.e., the job keeps the CPU until it blocks (say on an I/O device).

FCFS Scheduling Policy: Example

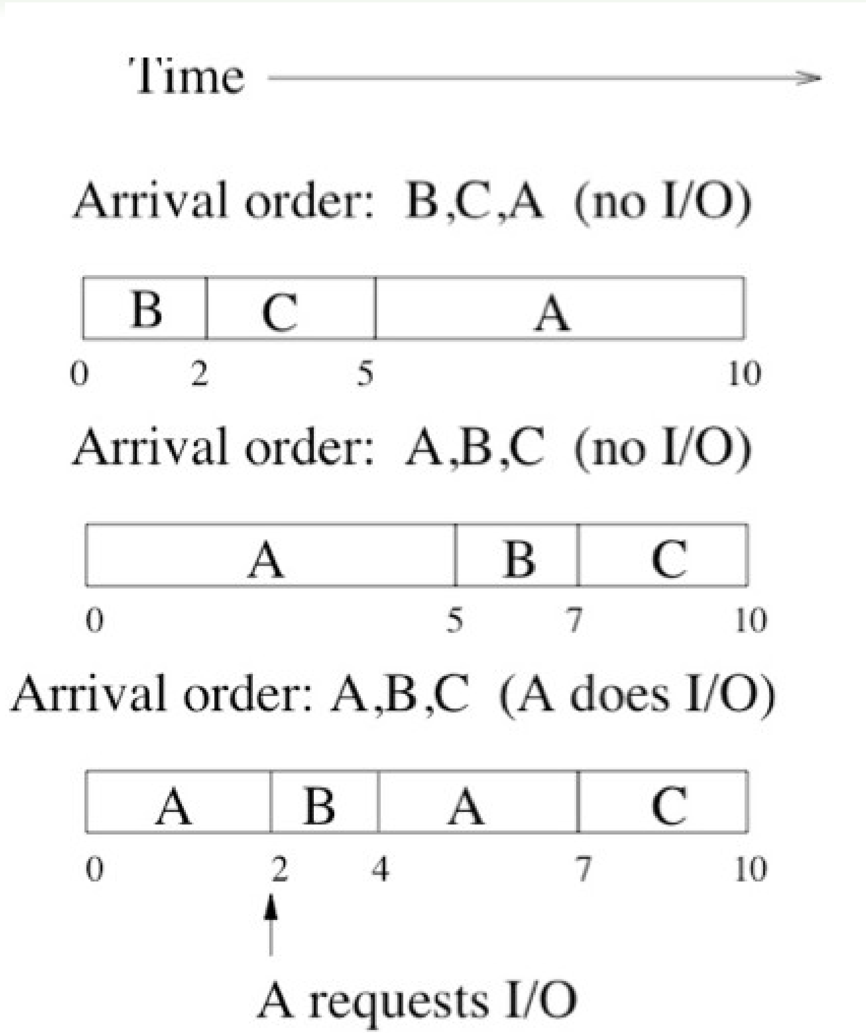

- If processes arrive 1 time unit apart, what is the average wait time in these three cases?

FCFS: pros and cons

- Advantage: simple
- Disadvantages:
  - average wait time is highly variable as short jobs may wait behind long jobs.
  - may lead to poor overlap of I/O and CPU since CPU-bound processes will force I/O bound processes to wait for the CPU, leaving the I/O devices idle

Round Robin Scheduling

- Variants of round robin are used in most time sharing systems
- Add a timer and use a preemptive policy.
- After each time slice, move the running thread to the back of the queue.

Round Robin Scheduling

- Selecting a time slice:
  - Too large - waiting time suffers, degenerates to FCFS if processes are never preempted.
  - Too small - throughput suffers because too much time is spent context switching.

  => Balance these tradeoffs by selecting a time slice where context switching is roughly 1% of the time slice.

  - typical time slice= 10-100 ms, context switch time= 0.1-1ms
- Advantage: It's fair; each job gets an equal shot at the CPU.
- Disadvantage: Average waiting time can be bad.

Round Robin Scheduling: Example 1

- 5 jobs, 100 seconds each, time slice 1 second, context switch time of 0

Job	Length	Completion Time		Wait Time	
		FCFS	Round Robin	FCFS	Round Robin
1	100				
2	100				
3	100				
4	100				
5	100				
Average					

Round Robin Scheduling: Example 1

- 5 jobs, 100 seconds each, time slice 1 second, context switch time of 0

Job	Length	Completion Time		Wait Time	
		FCFS	Round Robin	FCFS	Round Robin
1	100	100	496	0	396
2	100	200	497	100	397
3	100	300	498	200	398
4	100	400	499	300	399
5	100	500	500	400	400
Average		300	498	200	398

Round Robin Scheduling: Example 2

- 5 jobs, of length 50, 40, 30, 20, and 10 seconds each, time slice 1 second, context switch time of 0 seconds

Job	Length	Completion Time		Wait Time	
		FCFS	Round Robin	FCFS	Round Robin
1	50				
2	40				
3	30				
4	20				
5	10				
Average					

Round Robin Scheduling: Example 2

- 5 jobs, of length 50, 40, 30, 20, and 10 seconds each, time slice 1 second, context switch time of 0 seconds

Job	Length	Completion Time		Wait Time	
		FCFS	Round Robin	FCFS	Round Robin
1	50	50	150	0	100
2	40	90	140	50	100
3	30	120	120	90	90
4	20	140	90	120	70
5	10	150	50	140	40
Average		110	110	80	80

Scheduling Algorithms

- FCFS: First Come, First Served
- Round Robin: Use a time slice and preemption to alternate jobs.
- SJF: Shortest Job First
- Multilevel Feedback Queues: Round robin on each priority queue.
- Lottery Scheduling: Jobs get tickets and scheduler randomly picks winning ticket.

SJF/SRTF: Shortest Job First

- Schedule the job that has the least (expected) amount of work (CPU time) to do until its next I/O request or termination.
- Advantages:
  - Provably optimal with respect to minimizing the average waiting time
  - Works for preemptive and non-preemptive schedulers
  - Preemptive SJF is SRTF - shortest remaining time first

  => I/O bound jobs get priority over CPU bound jobs
- Disadvantages:
  - Impossible to predict the amount of CPU time a job has left
  - Long running CPU bound jobs can starve

SJF: Example

- 5 jobs, of length 50, 40, 30, 20, and 10 seconds each, time slice 1 second, context switch time of 0 seconds

Job	Length	**Completion Time**			**Wait Time**		
		FCFS	RR	SJF	FCFS	RR	SJF
1	50						
2	40						
3	30						
4	20						
5	10						
Average							

SJF: Example

- 5 jobs, of length 50, 40, 30, 20, and 10 seconds each, time slice 1 second, context switch time of 0 seconds

Job	Length	Completion Time			Wait Time		
		FCFS	RR	SJF	FCFS	RR	SJF
1	50	50	150	150	0	100	100
2	40	90	140	100	50	100	60
3	30	120	120	60	90	90	30
4	20	140	90	30	120	70	10
5	10	150	50	10	140	40	0
Average		110	110	70	80	80	40

Multilevel Feedback Queues (MLFQ)

- Multilevel feedback queues use past behavior to predict the future and assign job priorities

  => overcome the prediction problem in SJF
- If a process is I/O bound in the past, it is also likely to be I/O bound in the future (programs turn out not to be random.)
- To exploit this behavior, the scheduler can favor jobs that have used the least amount of CPU time, thus approximating SJF.
- This policy is adaptive because it relies on past behavior and changes in behavior result in changes to scheduling decisions.

Approximating SJF: Multilevel Feedback Queues

- Multiple queues with different priorities.
- Use Round Robin scheduling at each priority level, running the jobs in highest priority queue first.
- Once those finish, run jobs at the next highest priority queue, etc. (Can lead to starvation.)
- Round robin time slice increases exponentially at lower priorities.

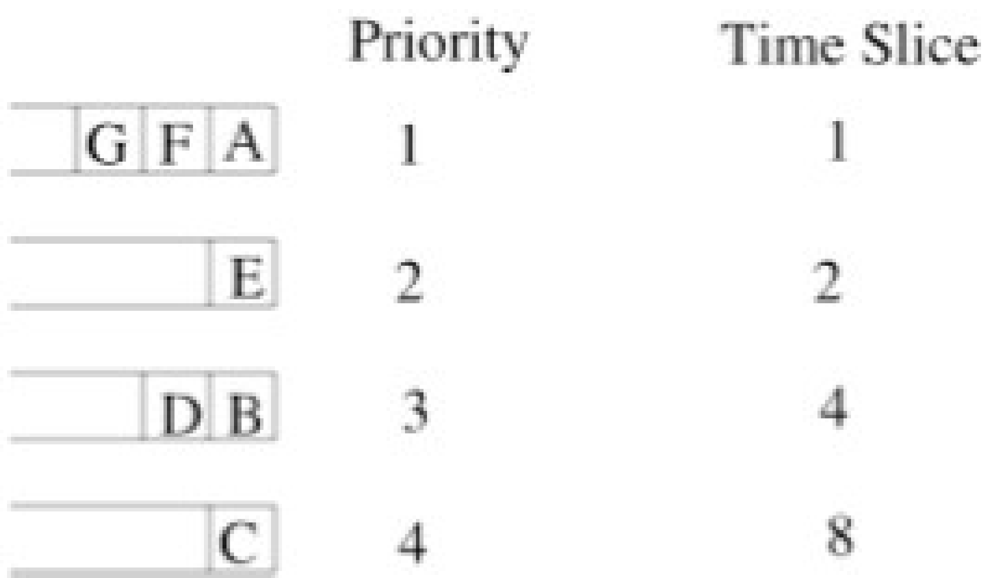

Adjusting Priorities in MLFQ

- Job starts in highest priority queue.
- If job's time slices expires, drop its priority one level.
- If job's time slices does not expire (the context switch comes from an I/O request instead), then increase its priority one level, up to the top priority level.

  => CPU bound jobs drop like a rock in priority and I/O bound jobs stay at a high priority.

Multilevel Feedback Queues: Example 1

- 3 jobs, of length 30, 20, and 10 seconds each, initial time slice 1 second, context switch time of 0 seconds, all CPU bound (no I/O), 3 queues

Job	Length	Completion Time		Wait Time	
		RR	MLFQ	RR	MLFQ
1	30				
2	20				
3	10				
Average					

Queue	Time Slice	Job
1	1	
2	2	
3	4	

Multilevel Feedback Queues: Example 1

- 3 jobs, of length 30, 20, and 10 seconds each, initial time slice 1 second, context switch time of 0 seconds, all CPU bound (no I/O), 3 queues

Job	Length	Completion Time		Wait Time	
		RR	MLFQ	RR	MLFQ
1	30	60	60	30	30
2	20	50	53	30	33
3	10	30	32	20	22
Average		46 2/3	48 1/3	26	28 1/3

Queue	Time Slice	Job
1	1	$1_1^1, 2_2^1, 3_3^1$
2	2	$1_5^3, 2_7^3, 3_9^3$
3	4	$1_{13}^7, 2_{17}^7, 3_{21}^7$ $1_{25}^{11}, 2_{29}^{11}, 3_{32}^{10}$ ...

Improving Fairness

- Since SJF is optimal, but unfair, any increase in fairness by giving long jobs a fraction of the CPU when shorter jobs are available will degrade average waiting time.
- Possible solutions:
  - Give each queue a fraction of the CPU time. This solution is only fair if there is an even distribution of jobs among queues.
  - Adjust the priority of jobs as they do not get serviced (Unix originally did this.) This ad hoc solution avoids starvation but average waiting time suffers when the system is overloaded because all the jobs end up with a high priority.

Lottery Scheduling

- Give every job some number of lottery tickets.
- On each time slice, randomly pick a winning ticket.
- On average, CPU time is proportional to the number of tickets given to each job.
- Assign tickets by giving the most to short running jobs, and fewer to long running jobs (approximating SJF). To avoid starvation, every job gets at least one ticket.
- Degrades gracefully as load changes. Adding or deleting a job affects all jobs proportionately, independent of the number of tickets a job has.

Summary of Scheduling Algorithms

- **FCFS:** Not fair, and average waiting time is poor.
- **Round Robin:** Fair, but average waiting time is poor.
- **SJF:** Not fair, but average waiting time is minimized assuming we can accurately predict the length of the next CPU burst. Starvation is possible.
- **Multilevel Queuing:** An implementation (approximation) of SJF.
- **Lottery Scheduling:** Fairer with a low average waiting time, but less predictable.