CPU Scheduling

COMP 3361: Operating Systems I Winter 2015 http://www.cs.du.edu/3361







- Selects from among the processes in the ready state and allocates the CPU to one of them
- CPU scheduling decisions may take place when a process
 - is created
 - terminates
 - switches from running to waiting state
 - > an interrupt occurs (I/O complete, timer, etc.)
- Operations performed to do process switching is pure overhead

Scheduling Types

- Non-preemptive: Once a process starts running, it is allowed to run until it blocks
 - after an interrupt is handled, CPU goes back to the process running prior to the interrupt
- Preemptive: A process may be forcibly removed from the CPU
 - after interrupt is handled, CPU may not go back to the process running prior to the interrupt

Invoking the Scheduler

Scenario:

- user types command to start a program
- the program starts running
- user is unable to type anything on the console until program finishes
- Why?
- We need to invoke the scheduler periodically
- Timer interrupts can be used to hand over control to the OS at regular intervals
 - configure timer hardware to generate an interrupt after fixed interval
 - timer interrupt handler invokes scheduler

Scheduling Algorithm Goals

Any system

fairness, policy enforcement, balance

Batch systems

- throughput, turnaround time, CPU utilization
- Interactive systems
 - response time, proportionality

Real-time systems

meeting deadlines, predictability





First-Come First-Served • If ready queue is P_2 , P_3 , P_1 P_1 P_2 P_3 3 30 6 • Waiting time for $P_1 = 6$; $P_2 = 0$; $P_3 = 3$ • Average waiting time = (6+0+3)/3 = 3

- much better than the previous case
- convoy effect is prevented
 - many I/O bound processes behind a CPU bound process
- Preemptive or non-preemptive?

- Associate with each process the length of its next CPU burst
- Use these lengths to first schedule the process with the shortest time
- SJF schedule gives the minimum average waiting time for a given set of processes
 - the difficulty is knowing the length of the next CPU burst

SJF Example

Process	Burst Time
P	6
P ₂	8
P ₃	7
P ₄	3

SJF Gantt chart

F) 4	P ₁		P ₃	P ₂		
0		3	9	1	6	24	

Average waiting time: (3+16+9+0)/4 = 7

12

Length of Next CPU Burst

Can only estimate the length

- Can be done by using the length of previous CPU bursts, using exponential averaging
 - t_n = actual length of n^{th} CPU burst
 - τ_{n+1} = predicted value for the next CPU burst
 - $\alpha, 0 \le \alpha \le 1$
 - Define: $\tau_{n+1} = \alpha t_n + (1 \alpha) \tau_n$

4 Example of Exponential Averaging α = 0 τ_{n+1} = τ_n; recent history does not count

- α = I
 - $\tau_{n+1} = t_n$; only the actual last CPU burst counts
- If we expand the formula, we get:

 $\begin{aligned} \tau_{n+1} &= \alpha \ t_n + (1 - \alpha) \alpha \ t_{n-1} + \dots \\ &+ (1 - \alpha)^j \alpha \ t_{n-j} + \dots \\ &+ (1 - \alpha)^{n+1} \tau_0 \end{aligned}$

Since both α and (I - α) are less than or equal to I, each successive term has less weight than its predecessor

Shortest Remaining Time First

- Preemptive shortest job first
 - always run the job that can finish the earliest

Process	Arrival Time	Burst Time		
P	0	8		
P ₂	I	4		
P ₃	2	9		
P ₄	3	5		





- Scheduler switches to next ready process in the queue
 - when the running process leaves the CPU voluntarily
 - when the timer hardware generates an interrupt

Performance

- long quantum \Rightarrow FIFO, not good for interactive environments
- Short quantum ⇒ quantum must be large with respect to context switch time, otherwise overhead is too high

RR Example

Process	Burst Time
PI	24
P ₂	3
P ₃	3

time quantum = 4

P ₁	P ₂	P ₃	P ₁				
) 4	L 7	7 1	0 1	4 1	8 2	2 2	6 30

• Higher average turnaround than SJF, but better response

CPU Scheduling

RR with CPU Shares

- Process P_1 , P_2 and P_3
- Assume that P₁ is guaranteed a 50% CPU share; the remaining is divided between the rest

Round Robin schedules

$$\begin{array}{|c|c|c|c|c|c|c|c|c|} P_1 & P_2 & P_1 & P_2 & P_1 & P_2 & P_1 & P_3 & P_1 & \dots \end{array}$$

- A priority number (integer) is associated with each process
- The CPU is allocated to the process with the highest priority
 - preemptive or non-preemptive?
 - SJF is a priority scheduler where priority is the predicted next CPU burst time
- Assign high priority to I/O bound processes Why?



Multiple Queues

Three queues

22

- > Q_0 : RR with time quantum 8 milliseconds
- Q_1 : RR with time quantum 16 milliseconds

• Q_2 : FCFS



Scheduling Threads in User Space



- I. Kernel scheduler picks process A to run
- Scheduler in thread run-time system decides the order in which process A's threads run in the time quantum
- 3. Kernel scheduler picks process B to run
- Scheduler in thread run-time system decides the order in which process B's threads run in the time quantum

likely order: $A_1, A_2, A_3, A_1, A_2, A_3, A_1, A_2, A_3$ unlikely order: $A_1, A_2, A_3, B_1, B_2, A_1, A_2, A_3, B_1$

Scheduling Threads in Kernel Space



- I. Kernel scheduler picks a thread from the thread table
- 2. When the time quantum ends, kernel scheduler picks another thread from the thread table
 - may belong to same process, or to a different one



References

 Chapter 2.4, Modern Operating Systems, A. Tanenbaum and H. Bos, 4th Edition.