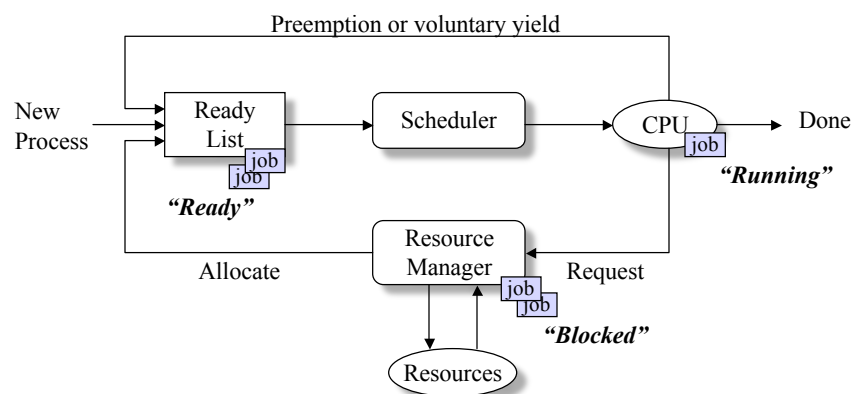


# Scheduling: the act of sharing



Copyright © 2004 Pearson Education, Inc.

## Model of Process Execution



## Thread leaves CPU

Slide 7-3

1. Completes execution
2. Requests a resource
3. Voluntarily release CPU
4. Involuntarily releases CPU (preempted)



Scheduling Policy - decides when thread leaves and which thread is selected to replace it

Scheduling mechanism - determines how the process manager carries out its duties

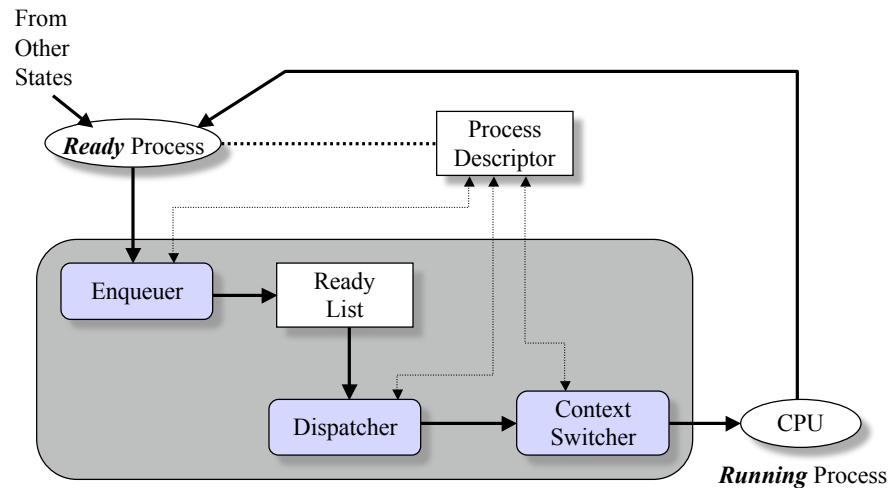
## Scheduling Mechanism

Slide 7-4

- Depends on hardware
  - Need a clock device
  - Rest of scheduler implemented in OS software
- 3 logical parts
  1. Enqueuer
  2. Dispatcher
  3. Context switcher

# The Scheduler

Slide 7-5



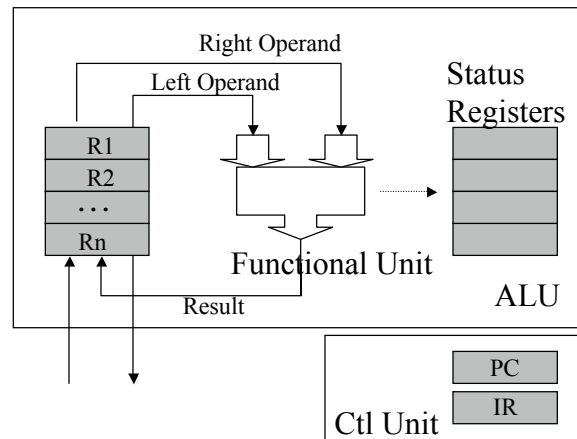
## Scheduling Mechanism

Slide 7-6

1. **Enqueuer**  
Places a pointer to thread descriptor into ready list (priority placement may happen then or later when thread is selected)
2. **Dispatcher**  
Selects one ready thread and allocates CPU to that thread by performing context switch
3. **Context switcher**  
Saves the contents of all CPU registers (PC, IR, condition status) for the thread being moved out and places the new thread in.


## Process/Thread Context

Slide 7-7



## Context Switching

Slide 7-8



## Context Switching

Slide 7-9

- Switching Cost
  - Over 32 different register - 32 or 64 but registers
  - With conventional load and stores  
(General registers + status registers)store ops X time/store

### 4 total context switches per multiplex

1. Save process 1's context
2. Load dispatcher's context
3. Save dispatcher's context
4. Load process 2's context

## Context Switching

Slide 7-10

- Switching Cost
  - 50 nanoseconds to store 1 unit of data
  - 32 bit bus / 50 nanoseconds for each register
  - 32 general registers and 8 status registers
    - $40 \times 50 \text{ nanoseconds} = 2 \text{ microseconds}$
  - 2 microseconds more to load another thread
  - Total time: over 4 microseconds (including dispatch in and out)
- 1 GHz processor - 2 nanosecond instructions then CPU could execute about 2,000 instructions while switch is taking place.

## Invoking the Scheduler

Slide 7-11

- Need a *mechanism* to call the scheduler
- Voluntary call
  - Process blocks itself
  - Calls the scheduler
- Involuntary call
  - External force (interrupt) blocks the process
  - Calls the scheduler

## Voluntary CPU Sharing

Slide 7-12

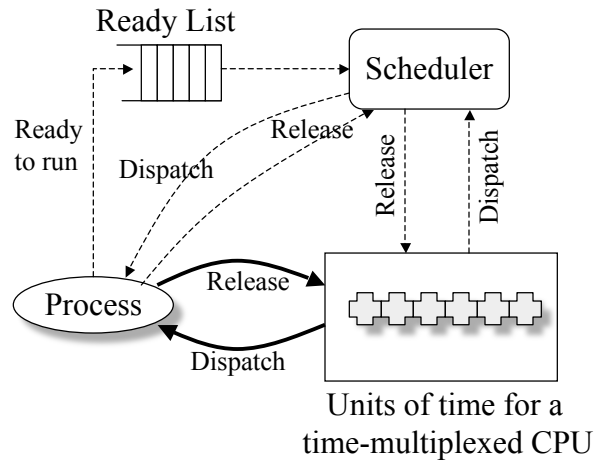
### Yield Instruction

```
yield( $p_i.pc$ ,  $p_j.pc$ ) {  
    memory[ $p_i.pc$ ] = PC;  
    PC = memory[ $p_j.pc$ ];  
}
```

- $p_i$  can be “automatically” determined from the processor status registers

```
yield(*,  $p_j.pc$ ) {  
    memory[ $p_i.pc$ ] = PC;  
    PC = memory[ $p_j.pc$ ];  
}
```

## Scheduler as CPU Resource Manager



## More on Yield

- $p_i$  and  $p_j$  can resume one another's execution

```
yield(*, p_j.pc);
...
yield(*, p_i.pc);
...
yield(*, p_j.pc);
...
```

- Suppose  $p_j$  is the scheduler:

```
// p_i yields to scheduler
yield(*, p_j.pc);
// scheduler chooses p_k
yield(*, p_k.pc);
// p_k yields to scheduler
yield(*, p_j.pc);
// scheduler chooses ...
```

## Voluntary Sharing

Slide 7-15

### Nonpreemptive Scheduler

- Scheduler using voluntary CPU sharing
- used in Xerox Alto PC
- used in earlier version of Mac OS

Problem:

What happens if some processes do not voluntarily yield?

## Voluntary Sharing

Slide 7-16

- Every process must periodically yield to the scheduler
- Relies on correct process behavior
  - Malicious
  - Accidental (infinite loop)
- Need a mechanism to override running process



## Involuntary CPU Sharing

Slide 7-17

- Interval timer
  - Device to produce a periodic interrupt
  - Programmable period

```
IntervalTimer() {  
    InterruptCount--;  
    if(InterruptCount <= 0) {  
        InterruptRequest = TRUE;  
        InterruptCount = K;  
    }  
    SetInterval(programmableValue) {  
        K = programmableValue;  
        InterruptCount = K;  
    }  
}
```

## Involuntary CPU Sharing (cont)

Slide 7-18

- Interval timer device
  - Serviced by a timer handler
    - Keeps an in-memory clock up-to-date
    - Invokes the scheduler

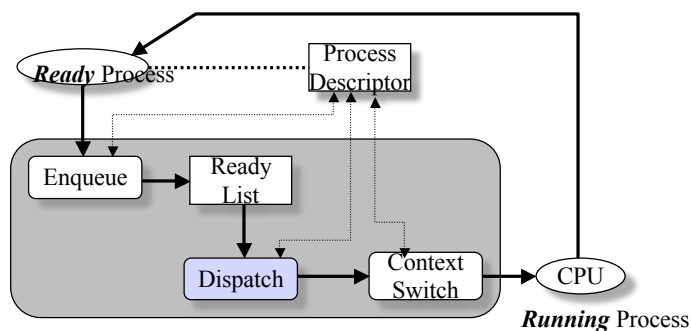
```
IntervalTimerHandler() {  
    Time++; // update the clock  
    TimeToSchedule--; // update alarm  
    if(TimeToSchedule <= 0) {  
        <invoke scheduler>;  
        TimeToSchedule = TimeSlice;  
    }  
}
```

# Contemporary Scheduling

- Involuntary CPU sharing – timer interrupts
  - Time quantum determined by interval timer – usually fixed size for every process using the system
  - Sometimes called the time slice length

## Preemptive Scheduler

## Choosing a Process to Run



- Mechanism never changes
- Strategy = policy the dispatcher uses to select a process from the ready list
- Different policies for different requirements

## Policy Considerations

Slide 7-21

- Policy can control/influence:
  - CPU utilization
  - Average time a process waits for service
  - Average amount of time to complete a job
- Could strive for any of:
  - Equitability (watch out for starvation)
  - Favor very short or long jobs
  - Meet priority requirements
  - Meet deadlines

## Optimal Scheduling

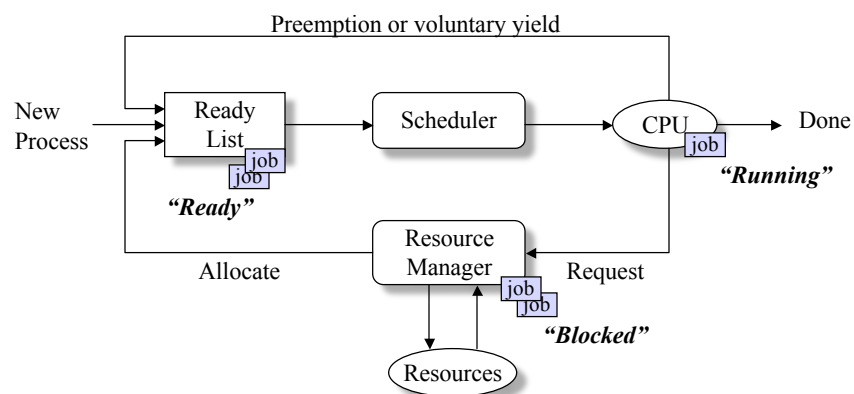
Slide 7-22

- Suppose the scheduler knows each process  $p_i$ 's service time,  $\tau(p_i)$  -- or it can estimate each  $\tau(p_i)$  :
- Policy can optimize on any criteria, e.g.,
  - CPU utilization
  - Waiting time
  - Deadline (i.e. real-time system)
- To find an optimal schedule:
  - Must have a finite, fixed # of  $p_i$
  - Know  $\tau(p_i)$  for each  $p_i$
  - Enumerate all schedules, then choose the best

## However ...

- The  $\tau(p_i)$  are almost certainly just estimates
- General algorithm to choose optimal schedule is  $O(n^2)$
- Other processes may arrive while these processes are being serviced
- Usually, optimal schedule is only a *theoretical benchmark* – scheduling policies try to *approximate* an optimal schedule

## Model of Process Execution



## Talking About Scheduling ...

Slide 7-25

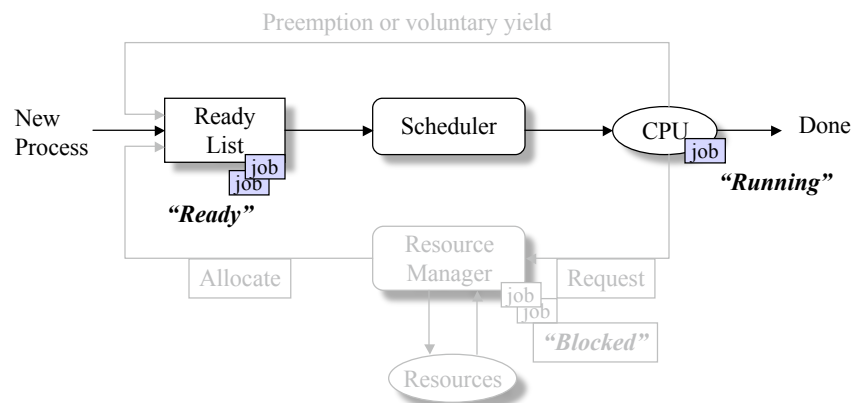
- Let  $P = \{p_i \mid 0 \leq i < n\}$  = set of processes
- Let  $S(p_i) \in \{\text{running, ready, blocked}\}$

### Common metrics for comparing scheduling strategies

1. Let  $\tau(p_i)$  = Time process needs to be in running state (the service time)
  2. Let  $W(p_i)$  = Time  $p_i$  is in ready state before first transition to running (wait time)
  3. Let  $T_{\text{TRnd}}(p_i)$  = Time from  $p_i$  first enter ready to last exit running (turnaround time)
- Batch Throughput rate = inverse of avg  $T_{\text{TRnd}}$
  - Timesharing response time =  $W(p_i)$


## Simplified Model

Slide 7-26




- Simplified, but still can provide analysis result
- Easy to analyze performance
- No issue of voluntary/involuntary CPU sharing

## Estimating CPU Utilization




Let  $\lambda$  = the average rate  
at which processes are  
placed in the Ready List,  
arrival rate

Let  $\mu$  = the average service rate  
 $\therefore 1/\mu$  = the average  $\tau(p_i)$



## Estimating CPU Utilization



Let  $\lambda$  = the average rate  
at which processes are  
placed in the Ready List,  
arrival rate


Let  $\mu$  = the average service rate  
 $\therefore 1/\mu$  = the average  $\tau(p_i)$

Let  $\rho$  = the fraction of the time that the CPU is expected to be busy  
 $\rho = \# p_i \text{ that arrive per unit time} * \text{avg time each spends on CPU}$   
 $\rho = \lambda * 1/\mu = \lambda/\mu$

- Notice must have  $\lambda < \mu$  (i.e.,  $\rho < 1$ )
- What if  $\rho$  approaches 1?

## Nonpreemptive Schedulers


Slide 7-29



- Try to use the simplified scheduling model
- Only consider running and ready states
- Ignores time in blocked state:
  - “New process created when it enters ready state”
  - “Process is destroyed when it enters blocked state”
  - Really just looking at “small phases” of a process

## First-Come-First-Served


Slide 7-30



$$T_{\text{TRnd}}(p_0) = \tau(p_0) = 350$$

$$W(p_0) = 0$$

## First-Come-First-Served




$$T_{\text{TRnd}}(p_0) = \tau(p_0) = 350$$

$$W(p_0) = 0$$

$$T_{\text{TRnd}}(p_1) = (\tau(p_1) + T_{\text{TRnd}}(p_0)) = 125 + 350 = 475$$

$$W(p_1) = T_{\text{TRnd}}(p_0) = 350$$

## First-Come-First-Served



$$T_{\text{TRnd}}(p_0) = \tau(p_0) = 350$$

$$W(p_0) = 0$$

$$T_{\text{TRnd}}(p_1) = (\tau(p_1) + T_{\text{TRnd}}(p_0)) = 125 + 350 = 475$$

$$W(p_1) = T_{\text{TRnd}}(p_0) = 350$$


$$T_{\text{TRnd}}(p_2) = (\tau(p_2) + T_{\text{TRnd}}(p_1)) = 475 + 475 = 950$$

$$W(p_2) = T_{\text{TRnd}}(p_1) = 475$$



## First-Come-First-Served

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75



$$T_{\text{TRnd}}(p_0) = \tau(p_0) = 350$$

$$W(p_0) = 0$$

$$T_{\text{TRnd}}(p_1) = (\tau(p_1) + T_{\text{TRnd}}(p_0)) = 125 + 350 = 475$$

$$W(p_1) = T_{\text{TRnd}}(p_0) = 350$$

$$T_{\text{TRnd}}(p_2) = (\tau(p_2) + T_{\text{TRnd}}(p_1)) = 475 + 475 = 950$$


$$W(p_2) = T_{\text{TRnd}}(p_1) = 475$$

$$T_{\text{TRnd}}(p_3) = (\tau(p_3) + T_{\text{TRnd}}(p_2)) = 250 + 950 = 1200$$

$$W(p_3) = T_{\text{TRnd}}(p_2) = 950$$

## First-Come-First-Served

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75



$$T_{\text{TRnd}}(p_0) = \tau(p_0) = 350$$

$$W(p_0) = 0$$

$$T_{\text{TRnd}}(p_1) = (\tau(p_1) + T_{\text{TRnd}}(p_0)) = 125 + 350 = 475$$

$$W(p_1) = T_{\text{TRnd}}(p_0) = 350$$

$$T_{\text{TRnd}}(p_2) = (\tau(p_2) + T_{\text{TRnd}}(p_1)) = 475 + 475 = 950$$

$$W(p_2) = T_{\text{TRnd}}(p_1) = 475$$

$$T_{\text{TRnd}}(p_3) = (\tau(p_3) + T_{\text{TRnd}}(p_2)) = 250 + 950 = 1200$$

$$W(p_3) = T_{\text{TRnd}}(p_2) = 950$$

$$T_{\text{TRnd}}(p_4) = (\tau(p_4) + T_{\text{TRnd}}(p_3)) = 75 + 1200 = 1275$$

$$W(p_4) = T_{\text{TRnd}}(p_3) = 1200$$

## FCFS Average Wait Time

Slide 7-35

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

- Easy to implement
- Ignores service time, etc
- Not a great performer

0	350	475	950	1200	1275
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$	

$$\begin{aligned}
 T_{\text{TRnd}}(p_0) &= \tau(p_0) = 350 & W(p_0) &= 0 \\
 T_{\text{TRnd}}(p_1) &= (\tau(p_1) + T_{\text{TRnd}}(p_0)) = 125 + 350 = 475 & W(p_1) &= T_{\text{TRnd}}(p_0) = 350 \\
 T_{\text{TRnd}}(p_2) &= (\tau(p_2) + T_{\text{TRnd}}(p_1)) = 475 + 475 = 950 & W(p_2) &= T_{\text{TRnd}}(p_1) = 475 \\
 T_{\text{TRnd}}(p_3) &= (\tau(p_3) + T_{\text{TRnd}}(p_2)) = 250 + 950 = 1200 & W(p_3) &= T_{\text{TRnd}}(p_2) = 950 \\
 T_{\text{TRnd}}(p_4) &= (\tau(p_4) + T_{\text{TRnd}}(p_3)) = 75 + 1200 = 1275 & W(p_4) &= T_{\text{TRnd}}(p_3) = 1200
 \end{aligned}$$

$$W_{\text{avg}} = (0 + 350 + 475 + 950 + 1200) / 5 = 2974 / 5 = 595$$

$$T_{\text{TRnd}} = (350 + 475 + 950 + 1200 + 1275) / 5 = 4250 / 5 = 850$$

## Predicting Wait Time in FCFS

Slide 7-36

- In FCFS, when a process arrives, all in ready list will be processed before this job
- Let  $\mu$  be the service rate
- Let  $L$  be the ready list length
- $W_{\text{avg}}(p) = L * 1/\mu + 0.5 * 1/\mu = L/\mu + 1/(2\mu)$

## Predicting Wait Time in FCFS

Slide 7-37

- Let's compare predicted wait with actual in earlier example:

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

$$\begin{aligned}\text{Average service time} &= (350 + 125 + 475 + 250) / 4 \\ &= 1200 / 4 \\ &= 300 \text{ time units}\end{aligned}$$

For the 4th process -  $L = 3$  (3 in ready and 1 in CPU)

$$\begin{aligned}\text{Estimated waiting time} &= 3 / (1/300) + 1/2(300) \\ &= 900 + 150 = 1050\end{aligned}$$

(COMPARED TO ACTUAL WAIT TIME OF 1200)

## Shortest Job Next

Slide 7-38

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

$\xrightarrow{\quad\quad\quad}$ 

$0 \downarrow 75$ $p_4$
----------------------------

$$T_{\text{TRnd}}(p_4) = \tau(p_4) = 75$$

$$W(p_4) = 0$$

## Shortest Job Next

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	75	200
$p_4$	$p_1$	

$$T_{\text{TRnd}}(p_1) = \tau(p_1) + \tau(p_4) = 125 + 75 = 200$$

$$W(p_1) = 75$$

$$T_{\text{TRnd}}(p_4) = \tau(p_4) = 75$$

$$W(p_4) = 0$$

## Shortest Job Next

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	75	200	450
$p_4$	$p_1$	$p_3$	

$$T_{\text{TRnd}}(p_1) = \tau(p_1) + \tau(p_4) = 125 + 75 = 200$$

$$W(p_1) = 75$$


$$T_{\text{TRnd}}(p_3) = \tau(p_3) + \tau(p_1) + \tau(p_4) = 250 + 125 + 75 = 450$$

$$W(p_3) = 200$$

$$T_{\text{TRnd}}(p_4) = \tau(p_4) = 75$$

$$W(p_4) = 0$$

## Shortest Job Next



$$T_{\text{TRnd}}(p_0) = \tau(p_0) + \tau(p_3) + \tau(p_1) + \tau(p_4) = 350 + 250 + 125 + 75 = 800$$

$$W(p_0) = 450$$

$$T_{\text{TRnd}}(p_1) = \tau(p_1) + \tau(p_4) = 125 + 75 = 200$$

$$W(p_1) = 75$$


$$T_{\text{TRnd}}(p_3) = \tau(p_3) + \tau(p_1) + \tau(p_4) = 250 + 125 + 75 = 450$$

$$W(p_3) = 200$$

$$T_{\text{TRnd}}(p_4) = \tau(p_4) = 75$$

$$W(p_4) = 0$$

## Shortest Job Next



$$T_{\text{TRnd}}(p_0) = \tau(p_0) + \tau(p_3) + \tau(p_1) + \tau(p_4) = 350 + 250 + 125 + 75 = 800$$

$$W(p_0) = 450$$

$$T_{\text{TRnd}}(p_1) = \tau(p_1) + \tau(p_4) = 125 + 75 = 200$$

$$W(p_1) = 75$$

$$T_{\text{TRnd}}(p_2) = \tau(p_2) + \tau(p_0) + \tau(p_3) + \tau(p_1) + \tau(p_4) = 475 + 350 + 250 + 125 + 75 = 1275$$

$$W(p_2) = 800$$

$$T_{\text{TRnd}}(p_3) = \tau(p_3) + \tau(p_1) + \tau(p_4) = 250 + 125 + 75 = 450$$

$$W(p_3) = 200$$

$$T_{\text{TRnd}}(p_4) = \tau(p_4) = 75$$

$$W(p_4) = 0$$

## Shortest Job Next

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

- Minimizes wait time
- May starve large jobs
- Must know service times

0	75	200	450	800	1275
$p_4$	$p_1$	$p_3$	$p_0$	$p_2$	

$$T_{\text{TRnd}}(p_0) = \tau(p_0) + \tau(p_3) + \tau(p_1) + \tau(p_4) = 350 + 250 + 125 + 75 = 800$$

$$W(p_0) = 450$$

$$T_{\text{TRnd}}(p_1) = \tau(p_1) + \tau(p_4) = 125 + 75 = 200$$

$$W(p_1) = 75$$

$$T_{\text{TRnd}}(p_2) = \tau(p_2) + \tau(p_0) + \tau(p_3) + \tau(p_1) + \tau(p_4) = 475 + 350 + 250 + 125 + 75 = 1275$$

$$W(p_2) = 800$$

$$T_{\text{TRnd}}(p_3) = \tau(p_3) + \tau(p_1) + \tau(p_4) = 250 + 125 + 75 = 450$$

$$W(p_3) = 200$$

$$T_{\text{TRnd}}(p_4) = \tau(p_4) = 75$$

$$W(p_4) = 0$$

$$W_{\text{avg}} = (450 + 75 + 800 + 200 + 0) / 5 = 1525 / 5 = 305$$

## Priority Scheduling

i	$\tau(p_i)$	Pri
0	350	5
1	125	2
2	475	3
3	250	1
4	75	4

- Reflects importance of external use
- May cause starvation
- Can address starvation with aging

0	250	375	850	925	1275
$p_3$	$p_1$	$p_2$	$p_4$	$p_0$	

$$T_{\text{TRnd}}(p_0) = \tau(p_0) + \tau(p_4) + \tau(p_2) + \tau(p_1) + \tau(p_3) = 350 + 75 + 475 + 125 + 250 = 1275$$

$$W(p_0) = 925$$

$$T_{\text{TRnd}}(p_1) = \tau(p_1) + \tau(p_3) = 125 + 250 = 375$$

$$W(p_1) = 250$$

$$T_{\text{TRnd}}(p_2) = \tau(p_2) + \tau(p_1) + \tau(p_3) = 475 + 125 + 250 = 850$$

$$W(p_2) = 375$$

$$T_{\text{TRnd}}(p_3) = \tau(p_3) = 250$$

$$W(p_3) = 0$$

$$T_{\text{TRnd}}(p_4) = \tau(p_4) + \tau(p_2) + \tau(p_1) + \tau(p_3) = 75 + 475 + 125 + 250 = 925$$


$$W(p_4) = 850$$

$$W_{\text{avg}} = (925 + 250 + 375 + 0 + 850) / 5 = 2400 / 5 = 480$$


## Deadline Scheduling

i	$\tau(p_i)$	Deadline
0	350	575
1	125	550
2	475	1050
3	250	(none)
4	75	200

- Allocates service by deadline
- May not be feasible



## Preemptive Schedulers



- Highest priority process is guaranteed to be running at all times
  - Or at least at the beginning of a time slice
- Dominant form of contemporary scheduling
- But complex to build & analyze

## Round Robin (TQ=50)

i	$\tau(p_i)$		
0	350		
1	125		
2	475		
3	250	0 50	
4	75	<table><tr><td><math>p_0</math></td></tr></table>	$p_0$
$p_0$			

$$W(p_0) = 0$$

## Round Robin (TQ=50)

i	$\tau(p_i)$			
0	350			
1	125			
2	475			
3	250	0 100		
4	75	<table><tr><td><math>p_0</math></td><td><math>p_1</math></td></tr></table>	$p_0$	$p_1$
$p_0$	$p_1$			

$$W(p_0) = 0$$

$$W(p_1) = 50$$



## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0		100	
$p_0$	$p_1$	$p_2$	

$$\begin{aligned} W(p_0) &= 0 \\ W(p_1) &= 50 \\ W(p_2) &= 100 \end{aligned}$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0		100		200
$p_0$	$p_1$	$p_2$	$p_3$	

$$\begin{aligned} W(p_0) &= 0 \\ W(p_1) &= 50 \\ W(p_2) &= 100 \\ W(p_3) &= 150 \end{aligned}$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	100	200		
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$

$$\begin{aligned} W(p_0) &= 0 \\ W(p_1) &= 50 \\ W(p_2) &= 100 \\ W(p_3) &= 150 \\ W(p_4) &= 200 \end{aligned}$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	100	200	300	
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$

$$\begin{aligned} W(p_0) &= 0 \\ W(p_1) &= 50 \\ W(p_2) &= 100 \\ W(p_3) &= 150 \\ W(p_4) &= 200 \end{aligned}$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	100	200	300	400	475
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$	$p_0$
$p_1$	$p_2$	$p_3$	$p_4$	$p_0$	$p_1$

$$T_{\text{TRnd}}(p_4) = 475$$

$$\begin{aligned} W(p_0) &= 0 \\ W(p_1) &= 50 \\ W(p_2) &= 100 \\ W(p_3) &= 150 \\ W(p_4) &= 200 \end{aligned}$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	100	200	300	400	475	550
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$	$p_0$	$p_1$
$p_1$	$p_2$	$p_3$	$p_4$	$p_0$	$p_1$	$p_2$

$$T_{\text{TRnd}}(p_1) = 550$$

$$T_{\text{TRnd}}(p_4) = 475$$

$$\begin{aligned} W(p_0) &= 0 \\ W(p_1) &= 50 \\ W(p_2) &= 100 \\ W(p_3) &= 150 \\ W(p_4) &= 200 \end{aligned}$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	100	200	300	400	475	550	650
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$	$p_0$	$p_1$	$p_2$
$p_3$	$p_0$	$p_1$	$p_2$	$p_3$	$p_0$	$p_1$	$p_2$
$p_3$	$p_0$	$p_1$	$p_2$	$p_3$	$p_0$	$p_1$	$p_2$

$$T_{\text{TRnd}}(p_1) = 550$$

$$T_{\text{TRnd}}(p_3) = 950$$

$$T_{\text{TRnd}}(p_4) = 475$$

$$W(p_0) = 0$$

$$W(p_1) = 50$$

$$W(p_2) = 100$$

$$W(p_3) = 150$$

$$W(p_4) = 200$$

## Round Robin (TQ=50)

i	$\tau(p_i)$
0	350
1	125
2	475
3	250
4	75

0	100	200	300	400	475	550	650
$p_0$	$p_1$	$p_2$	$p_3$	$p_4$	$p_0$	$p_1$	$p_2$
$p_3$	$p_0$	$p_1$	$p_2$	$p_3$	$p_0$	$p_1$	$p_2$
$p_3$	$p_0$	$p_1$	$p_2$	$p_3$	$p_0$	$p_1$	$p_2$

$$T_{\text{TRnd}}(p_0) = 1100$$

$$T_{\text{TRnd}}(p_1) = 550$$

$$T_{\text{TRnd}}(p_3) = 950$$

$$T_{\text{TRnd}}(p_4) = 475$$

$$W(p_0) = 0$$

$$W(p_1) = 50$$

$$W(p_2) = 100$$

$$W(p_3) = 150$$

$$W(p_4) = 200$$

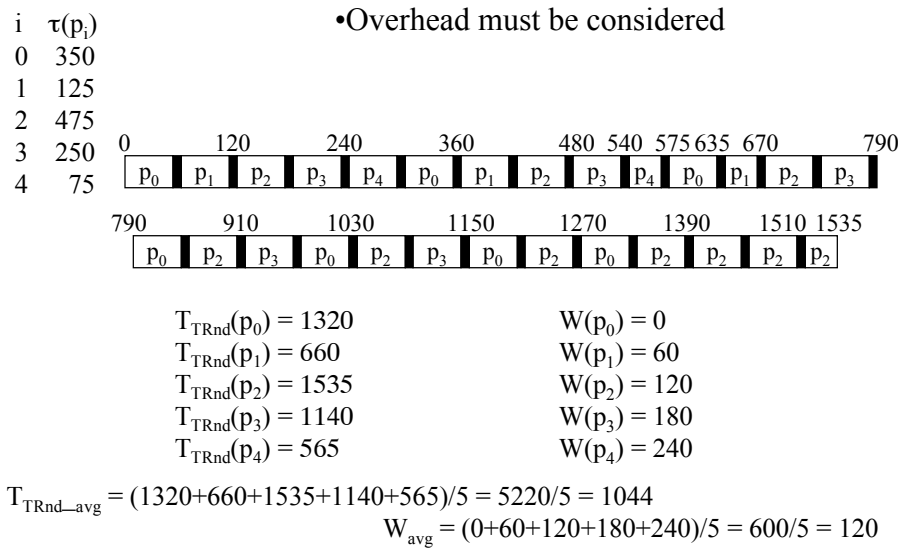
$T_{\text{TRnd}}(p_0) = 1100$	$W(p_0) = 0$
$T_{\text{TRnd}}(p_1) = 550$	$W(p_1) = 50$
$T_{\text{TRnd}}(p_2) = 1275$	$W(p_2) = 100$
$T_{\text{TRnd}}(p_3) = 950$	$W(p_3) = 150$
$T_{\text{TRnd}}(p_4) = 475$	$W(p_4) = 200$

- Equitable
- Most widely-used
- Fits naturally with interval timer

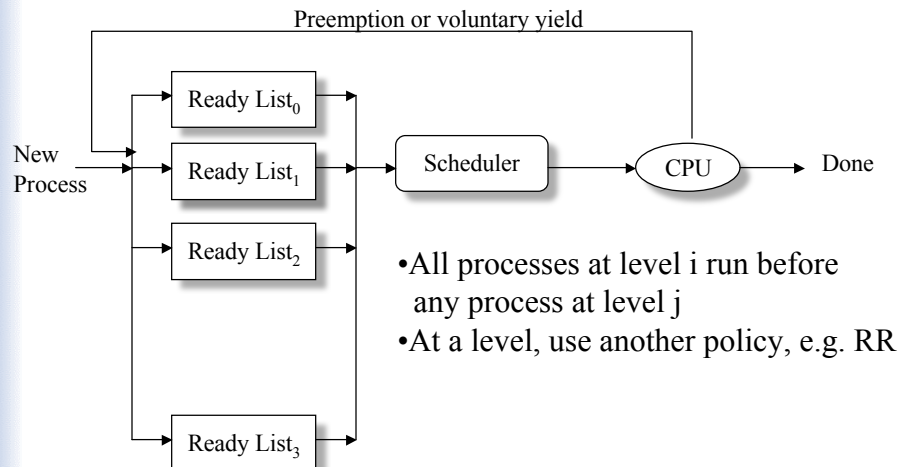
$T_{\text{TRnd}}(p_0) = 1100$	$W(p_0) = 0$
$T_{\text{TRnd}}(p_1) = 550$	$W(p_1) = 50$
$T_{\text{TRnd}}(p_2) = 1275$	$W(p_2) = 100$
$T_{\text{TRnd}}(p_3) = 950$	$W(p_3) = 150$
$T_{\text{TRnd}}(p_4) = 475$	$W(p_4) = 200$

$$W_{\text{avg}} = (0+50+100+150+200)/5 = 500/5 = 100$$

## RR with Overhead=10 (TQ=50)



## Multi-Level Queues



## Contemporary Scheduling

- Involuntary CPU sharing -- timer interrupts
  - Time quantum determined by interval timer -- usually fixed for every process using the system
  - Sometimes called the time slice length
- Priority-based process (job) selection
  - Select the highest priority process
  - Priority reflects policy
- With preemption
- Usually a variant of Multi-Level Queues

## BSD 4.4 Scheduling

- Involuntary CPU Sharing
- Preemptive algorithms
- 32 Multi-Level Queues
  - Queues 0-7 are reserved for system functions
  - Queues 8-31 are for user space functions
  - `nice` influences (but does not dictate) queue level


## Windows NT/2K Scheduling

Slide 7-63


- Involuntary CPU Sharing across threads
- Preemptive algorithms
- 32 Multi-Level Queues
  - Highest 16 levels are “real-time”
  - Next lower 15 are for system/user threads
    - Range determined by process base priority
  - Lowest level is for the idle thread

## Bank Teller Simulation

Slide 7-64



Tellers at the Bank



Model of Tellers at the Bank



## Simulation Kernel Loop

Slide 7-65

```
simulated_time = 0;
while (true) {
    event = select_next_event();
    if (event->time > simulated_time)
        simulated_time = event->time;
    evaluate(event->function, ...);
}
```

## Simulation Kernel Loop(2)

Slide 7-66

```
void runKernel(int quitTime)
{
    Event *thisEvent;

    // Stop by running to elapsed time, or by causing quit execute
    if(quitTime <= 0) quitTime = 9999999;
    simTime = 0;
    while(simTime < quitTime) {
        // Get the next event
        if(eventList == NIL) {
            // No more events to process
            break;
        }
        thisEvent = eventList;
        eventList = thisEvent->next;
        simTime = thisEvent->getTime(); // Set the time
        // Execute this event
        thisEvent->fire();
        delete(thisEvent);
    };
}
```