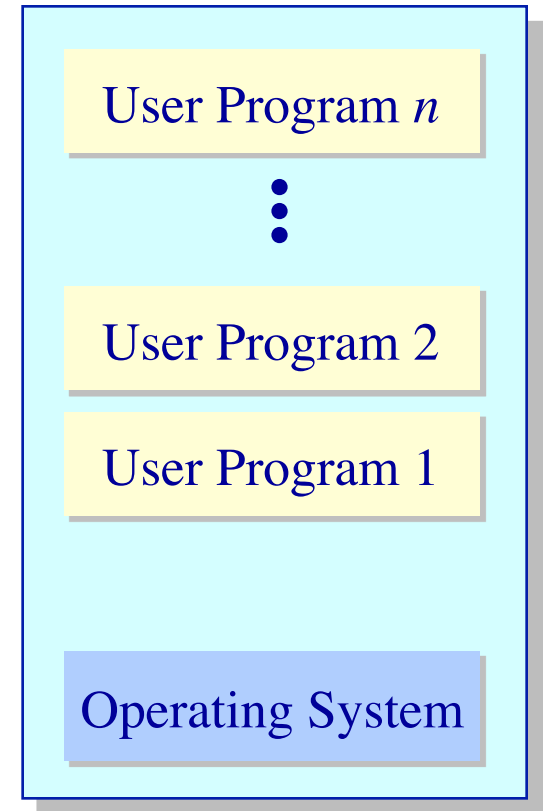


# *Page Replacement Algorithms*

# Virtual Memory Management

## Fundamental issues : A Recap

- ◆ Key concept: Demand paging
  - Load pages into memory only when a page fault occurs
- ◆ Issues:
  - Placement strategies
    - ❖ Place pages anywhere – no placement policy required
  - Replacement strategies
    - ❖ What to do when there exist more jobs than can fit in memory
  - Load control strategies
    - ❖ Determining how many jobs can be in memory at one time



Memory

# Page Replacement Algorithms

## Concept

- ◆ Typically  $\sum_i VAS_i \gg \text{Physical Memory}$
- ◆ With demand paging, physical memory fills quickly
- ◆ When a process faults & memory is full, some page must be swapped out
  - Handling a page fault now requires **2** disk accesses not 1!

Which page should be replaced?

*Local replacement* — Replace a page of the faulting process

*Global replacement* — Possibly replace the page of another process

# Page Replacement Algorithms

## Evaluation methodology

- ◆ Record a *trace* of the pages accessed by a process
  - Example: (Virtual page, offset) address trace...  
(3,0), (1,9), (4,1), (2,1), (5,3), (2,0), (1,9), (2,4), (3,1), (4,8)
  - generates page trace  
3, 1, 4, 2, 5, 2, 1, 2, 3, 4 (represented as *c, a, d, b, e, b, a, b, c, d*)
- ◆ Hardware can tell OS when a new page is loaded into the TLB
  - Set a used bit in the page table entry
  - Increment or shift a register

Simulate the behavior of a page replacement algorithm on the trace and record the number of page faults generated

*fewer faults*  *better performance*

# Optimal Page Replacement

## Clairvoyant replacement

- ◆ Replace the page that won't be needed for the longest time in the future

Initial allocation

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>									
1	<i>b</i>										
2	<i>c</i>										
3	<i>d</i>										
Faults											
Time page needed next											



# Optimal Page Replacement

## Clairvoyant replacement



- ◆ Replace the page that won't be needed for the longest time in the future

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>d</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>
Faults						•					•
Time page needed next					<i>a</i> = 7 <i>b</i> = 6 <i>c</i> = 9 <i>d</i> = 10					<i>a</i> = 15 <i>b</i> = 11 <i>c</i> = 13 <i>d</i> = 14	

# Local Page Replacement

## FIFO replacement

- ◆ Simple to implement
  - A single pointer suffices
- ◆ Performance with 4 page frames:




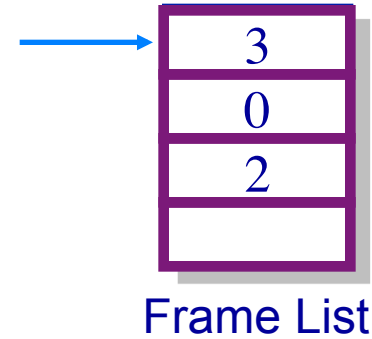
Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>									
	1	<i>b</i>									
	2	<i>c</i>									
	3	<i>d</i>									
Faults											



# Local Page Replacement

## FIFO replacement

- ◆ Simple to implement
  - A single pointer suffices
- ◆ Performance with 4 page frames:



Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>d</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>b</i>	<i>b</i>	<i>b</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
Faults						•		•	•	•	•



# Least Recently Used Page Replacement

Use the recent past as a predictor of the near future

- ◆ Replace the page that hasn't been referenced for the longest time

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>									
	1	<i>b</i>									
	2	<i>c</i>									
	3	<i>d</i>									
Faults											
Time page last used											



# Least Recently Used Page Replacement

Use the recent past as a predictor of the near future


- ◆ Replace the page that hasn't been referenced for the longest time

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>d</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
Faults						•				•	•
Time page last used				<i>a</i> = 2	<i>b</i> = 4	<i>c</i> = 1	<i>d</i> = 3		<i>a</i> = 7	<i>b</i> = 8	<i>c</i> = 9

# Least Recently Used Page Replacement Implementation

- ◆ Maintain a “stack” of recently used pages


Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>d</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
Faults						•				•	•



# Least Recently Used Page Replacement Implementation

- ◆ Maintain a “stack” of recently used pages

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>d</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
Faults						•				•	•




- ◆ What is the goal of a page replacement algorithm?
  - A. Make life easier for OS implementer
  - B. Reduce the number of page faults
  - C. Reduce the penalty for page faults when they occur
  - D. Minimize CPU time of algorithm

# Approximate LRU Page Replacement

## The *Clock* algorithm

- ◆ Maintain a circular list of pages resident in memory
  - Use a *clock* (or *used/referenced*) bit to track how often a page is accessed
  - The bit is set whenever a page is referenced
- ◆ Clock hand sweeps over pages looking for one with *used* bit = 0
  - Replace pages that haven't been referenced for one complete revolution of the clock



```
func Clock_Replacement  
begin  
  while (victim page not found) do  
    if (used bit for current page = 0) then  
      replace current page  
    else  
      reset used bit  
    end if  
    advance clock pointer  
  end while  
end Clock_Replacement
```

# Clock Page Replacement

## Example

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>						
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>						
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>						
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>						
Faults											

Page table entries  
for resident pages:

1	<i>a</i>
1	<i>b</i>
1	<i>c</i>
1	<i>d</i>









# Clock Page Replacement

## Example

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	<i>e</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>d</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
Faults						•		•		•	•

Page table entries  
for resident pages:

1	<i>a</i>
1	<i>b</i>
1	<i>c</i>
1	<i>d</i>

1	<i>e</i>
0	<i>b</i>
0	<i>c</i>
0	<i>d</i>

1	<i>e</i>
1	<i>b</i>
0	<i>c</i>
0	<i>d</i>

1	<i>e</i>
0	<i>b</i>
1	<i>a</i>
0	<i>d</i>

1	<i>e</i>
1	<i>b</i>
1	<i>a</i>
0	<i>d</i>

1	<i>e</i>
1	<i>b</i>
1	<i>a</i>
1	<i>c</i>

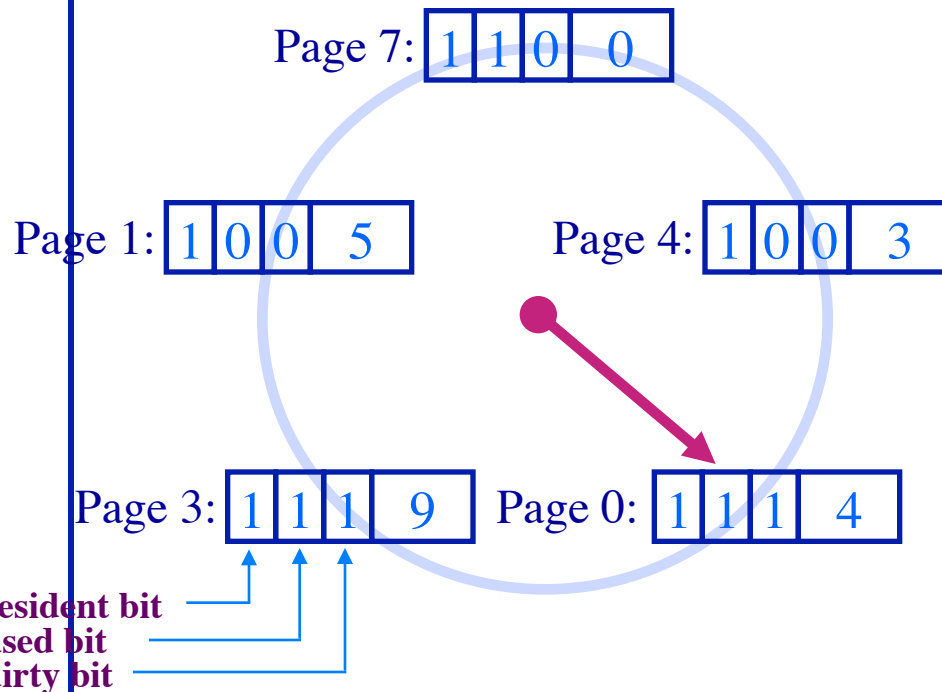
1	<i>d</i>
0	<i>b</i>
0	<i>a</i>
0	<i>c</i>



# Optimizing Approximate LRU Replacement

## The Second Chance algorithm

- ◆ There is a significant cost to replacing “dirty” pages
  - Why?
    - ❖ Must write back contents to disk before freeing!
- ◆ Modify the Clock algorithm to allow dirty pages to always survive one sweep of the clock hand
  - Use both the *dirty bit* and the *used bit* to drive replacement



### Second Chance Algorithm

Before clock sweep

<i>used</i>	<i>dirty</i>
0	0
0	1
1	0
1	1

After clock sweep

<i>used</i>	<i>dirty</i>
	<i>replace page</i>
0	0
0	0
0	1

# The Second Chance Algorithm

## Example

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a<sup>w</sup></i>	<i>d</i>	<i>b<sup>w</sup></i>	<i>e</i>	<i>b</i>	<i>a<sup>w</sup></i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>					
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>					
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>					
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>					

Faults

Page table  
entries  
for resident  
pages:

10	<i>a</i>
10	<i>b</i>
10	<i>c</i>
10	<i>d</i>









# The Second Chance Algorithm

## Example

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>a<sup>w</sup></i>	<i>d</i>	<i>b<sup>w</sup></i>	<i>e</i>	<i>b</i>	<i>a<sup>w</sup></i>	<i>b</i>	<i>c</i>	<i>d</i>
Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>d</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>	<i>e</i>
	3	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
Faults						•				•	•

Page table entries for resident pages:

10	<i>a</i>
10	<i>b</i>
10	<i>c</i>
10	<i>d</i>

11	<i>a</i>
11	<i>b</i>
10	<i>c</i>
10	<i>d</i>

00	<i>a</i> *
00	<i>b</i> *
10	<i>e</i>
00	<i>d</i>

00	<i>a</i>
10	<i>b</i>
10	<i>e</i>
00	<i>d</i>

11	<i>a</i>
10	<i>b</i>
10	<i>e</i>
00	<i>d</i>

11	<i>a</i>
10	<i>b</i>
10	<i>e</i>
10	<i>c</i>

00	<i>a</i> *
10	<i>d</i>
00	<i>e</i>
00	<i>c</i>

# The Problem With Local Page Replacement

How much memory do we allocate to a process?

Time	0	1	2	3	4	5	6	7	8	9	10	11	12
Requests		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>

Page Frames	0	<i>a</i>											
	1	<i>b</i>											
	2	<i>c</i>											
Faults													

Page Frames	0	<i>a</i>											
	1	<i>b</i>											
	2	<i>c</i>											
	3	-											
Faults													

# The Problem With Local Page Replacement

How much memory do we allocate to a process?

Time	0	1	2	3	4	5	6	7	8	9	10	11	12
Requests		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>

Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>b</i>	<i>b</i>	<i>b</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>d</i>
	Faults					•	•	•	•	•	•	•	•	•

Page Frames	0	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
	1	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
	2	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>
	3	-				<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>d</i>
Faults					•								

# Page Replacement Algorithms

## Performance

- ◆ Local page replacement
  - LRU — Ages pages based on when they were last used
  - FIFO — Ages pages based on when they're brought into memory
- ◆ Towards global page replacement ... with variable number of page frames allocated to processes

### The principle of locality

- 90% of the execution of a program is sequential
- Most iterative constructs consist of a relatively small number of instructions
- When processing large data structures, the dominant cost is sequential processing on individual structure elements
- Temporal vs. physical locality

# Optimal Page Replacement

For processes with a variable number of frames

- ◆ *VMIN* — Replace a page that is not referenced in the *next*  $\tau$  accesses
- ◆ Example:  $\tau = 4$

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>c</i>	<i>d</i>	<i>b</i>	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>	<i>a</i>	<i>d</i>
Pages in Memory	Page <i>a</i>	• <i>t=0</i>									
	Page <i>b</i>	-									
	Page <i>c</i>	-									
	Page <i>d</i>	• <i>t=-1</i>									
	Page <i>e</i>	-									
Faults											



# Optimal Page Replacement

For processes with a variable number of frames

- ◆ *VMIN* — Replace a page that is not referenced in the *next*  $\tau$  accesses
- ◆ Example:  $\tau = 4$

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>c</i>	<i>d</i>	<i>b</i>	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>	<i>a</i>	<i>d</i>
Pages in Memory	Page <i>a</i>	• <i>t=0</i>	-	-	-	-	-	-	-	<i>F</i>	-
	Page <i>b</i>	-	-	-	<i>F</i>	-	-	-	-	-	-
	Page <i>c</i>	-	<i>F</i>	•	•	•	•	•	•	-	-
	Page <i>d</i>	• <i>t=-1</i>	•	•	•	-	-	-	-	-	<i>F</i>
	Page <i>e</i>	-	-	-	-	-	<i>F</i>	•	•	-	-
Faults		•			•		•			•	•



# Explicitly Using Locality

## The working set model of page replacement

- ◆ Assume recently referenced pages are likely to be referenced again soon...
- ◆ ... and *only* keep those pages recently referenced in memory (called *the working set*)
  - Thus pages may be removed even when no page fault occurs
  - The number of frames allocated to a process will vary over time
- ◆ A process is allowed to execute only if its working set fits into memory
  - The working set model performs implicit load control

# Working Set Page Replacement Implementation

- ◆ Keep track of the last  $\tau$  references
  - The pages referenced during the last  $\tau$  memory accesses are the working set
  - $\tau$  is called the *window size*
- ◆ Example: Working set computation,  $\tau = 4$  references:

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>c</i>	<i>d</i>	<i>b</i>	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>	<i>a</i>	<i>d</i>
Pages in Memory	Page <i>a</i>	• <i>t</i> = 0									
	Page <i>b</i>	-									
	Page <i>c</i>	-									
	Page <i>d</i>	• <i>t</i> = -1									
	Page <i>e</i>	• <i>t</i> = -2									
Faults											



# Working Set Page Replacement Implementation

- ◆ Keep track of the last  $\tau$  references
  - The pages referenced during the last  $\tau$  memory accesses are the working set
  - $\tau$  is called the *window size*
- ◆ Example: Working set computation,  $\tau = 4$  references:
  - What if  $\tau$  is too small? too large?

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>c</i>	<i>d</i>	<i>b</i>	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>	<i>a</i>	<i>d</i>
Pages in Memory	Page <i>a</i>	$\bullet_{t=0}$	$\bullet$	$\bullet$	$\bullet$	-	-	-	-	$\circ F$	$\bullet$
	Page <i>b</i>	-	-	-	$\circ F$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	-	-
	Page <i>c</i>	-	$\circ F$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	$\bullet$
	Page <i>d</i>	$\bullet_{t=-1}$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	$\bullet$	-	-	-	$\circ F$
	Page <i>e</i>	$\bullet_{t=-2}$	$\bullet$	-	-	-	-	$\circ F$	$\bullet$	$\bullet$	$\bullet$
Faults		$\bullet$			$\bullet$		$\bullet$			$\bullet$	$\bullet$

# Page-Fault-Frequency Page Replacement

## An alternate working set computation

- ◆ Explicitly attempt to minimize page faults
  - When page fault frequency is high — *increase working set*
  - When page fault frequency is low — *decrease working set*

### Algorithm:

Keep track of the rate at which faults occur

When a fault occurs, compute the time since the last page fault

Record the time,  $t_{last}$ , of the last page fault

If the time between page faults is “large” then reduce the working set

If  $t_{current} - t_{last} > \tau$ , then remove from memory all pages not referenced in  $[t_{last}, t_{current}]$

If the time between page faults is “small” then increase working set

If  $t_{current} - t_{last} \leq \tau$ , then add faulting page to the working set

# Page-Fault-Frequency Page Replacement

## Example, window size = 2

- ◆ If  $t_{current} - t_{last} > 2$ , remove pages not referenced in  $[t_{last}, t_{current}]$  from the working set
- ◆ If  $t_{current} - t_{last} \leq 2$ , just add faulting page to the working set

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>c</i>	<i>d</i>	<i>b</i>	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>	<i>a</i>	<i>d</i>
Pages in Memory	Page <i>a</i>	•									
	Page <i>b</i>	-									
	Page <i>c</i>	-									
	Page <i>d</i>	•									
	Page <i>e</i>	•									
Faults											
$t_{cur} - t_{last}$											



# Page-Fault-Frequency Page Replacement

## Example, window size = 2

- ◆ If  $t_{current} - t_{last} > 2$ , remove pages not referenced in  $[t_{last}, t_{current}]$  from the working set
- ◆ If  $t_{current} - t_{last} \leq 2$ , just add faulting page to the working set

Time	0	1	2	3	4	5	6	7	8	9	10
Requests		<i>c</i>	<i>c</i>	<i>d</i>	<i>b</i>	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>	<i>a</i>	<i>d</i>
Pages in Memory	Page <i>a</i>	•	•	•	-	-	-	-	-	<i>F</i>	•
	Page <i>b</i>	-	-	-	<i>F</i>	•	•	•	•	-	-
	Page <i>c</i>	-	<i>F</i>	•	•	•	•	•	•	•	•
	Page <i>d</i>	•	•	•	•	•	•	•	•	-	<i>F</i>
	Page <i>e</i>	•	•	•	•	-	-	<i>F</i>	•	•	•
Faults		•			•		•			•	•
$t_{cur} - t_{last}$		1			3		2			3	1

# Load Control

## Fundamental tradeoff

- ◆ High multiprogramming level

- $MPL_{max} = \frac{\text{number of page frames}}{\text{minimum number of frames required for a process to execute}}$

- ◆ Low paging overhead

- $MPL_{min} = 1$  process


- ◆ Issues

- What criterion should be used to determine when to increase or decrease the  $MPL$ ?
  - Which task should be swapped out if the  $MPL$  must be reduced?

# Load Control

## How *not* to do it: Base load control on CPU utilization

- ◆ Assume memory is nearly full
- ◆ A chain of page faults occur
  - A queue of processes forms at the paging device
- ◆ CPU utilization falls
- ◆ Operating system increases *MPL*
  - New processes fault, taking memory away from existing processes
- ◆ CPU utilization goes to 0, the OS increases the *MPL* further...




System is *thrashing* — spending all of its time paging



# Load Control

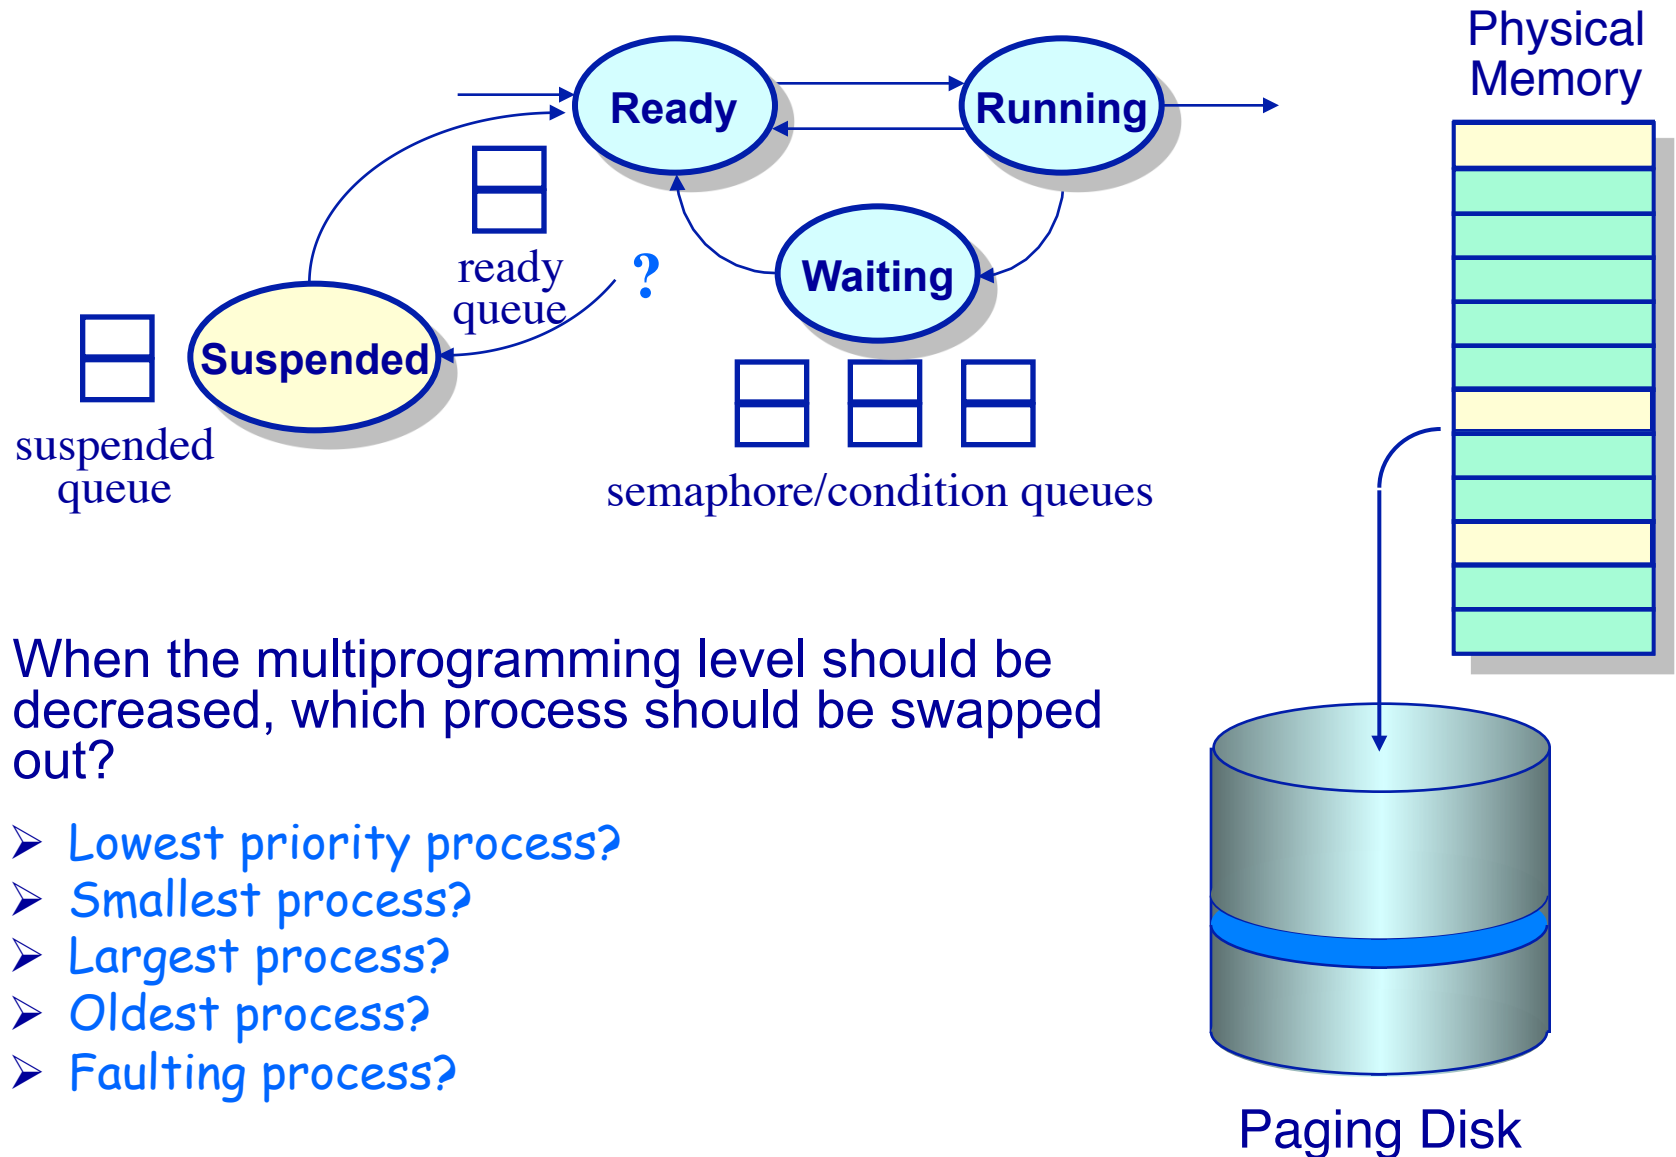
## Thrashing

- ◆ Thrashing can be ameliorated by *local* page replacement
- ◆ Better criteria for load control: Adjust MPL so that:
  - *mean time between page faults (MTBF) = page fault service time (PFST)*
  - $\sum WS_i = \text{size of memory}$



# Load Control

## Thrashing



- ◆ When the multiprogramming level should be decreased, which process should be swapped out?
  - Lowest priority process?
  - Smallest process?
  - Largest process?
  - Oldest process?
  - Faulting process?