OPERATING SYSTEMS

VIRTUAL MEMORY

Jerry Breecher

VIRTUAL MEMORY

WHY VIRTUAL MEMORY?

- We've previously required the entire logical space of the process to be in memory before the process could run. We will now look at alternatives to this.

- Most code/data isn't needed at any instant, or even within a finite time - we can bring it in only as needed.

VIRTUES

- Gives a higher level of multiprogramming

- The program size isn't constrained (thus the term 'virtual memory'). Virtual memory allows very large logical address spaces.

- Swap sizes smaller.

VIRTUAL MEMORY

Definitions

Virtual memory

The conceptual separation of user logical memory from physical memory. Thus we can have large virtual memory on a small physical memory.

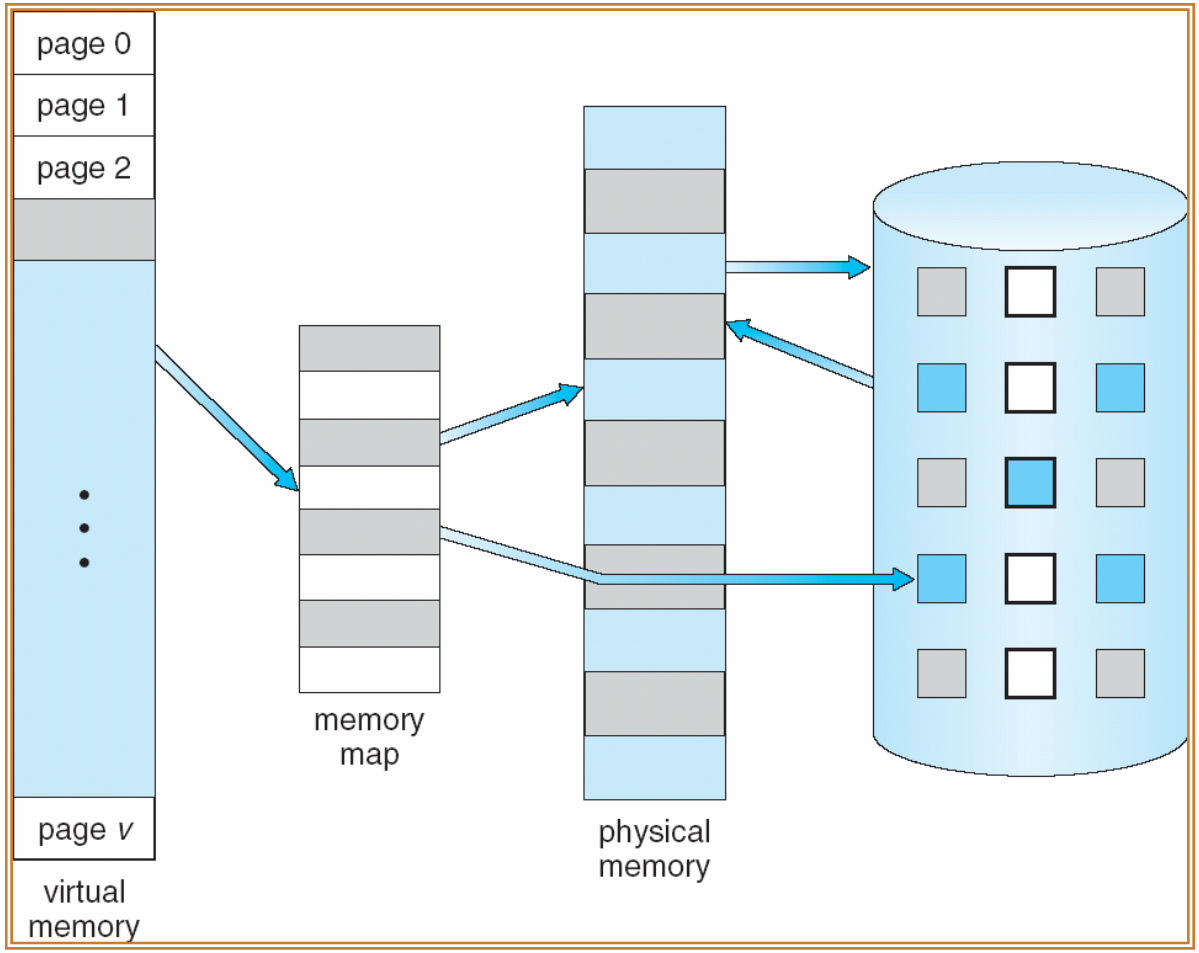

VIRTUAL MEMORY

Definitions

Demand paging When a page is touched, bring it from secondary to main memory.

Overlays Laying of code data on the same logical addresses - this is the reuse of logical memory. Useful when the program is in phases or when logical address space is small.

Dynamic loading A routine is loaded only when it's called.

VIRTUAL MEMORY

Demand Paging

When a page is referenced, either as code execution or data access, and that page isn't in memory, then get the page from disk and re-execute the statement.

Here's migration between memory and disk.

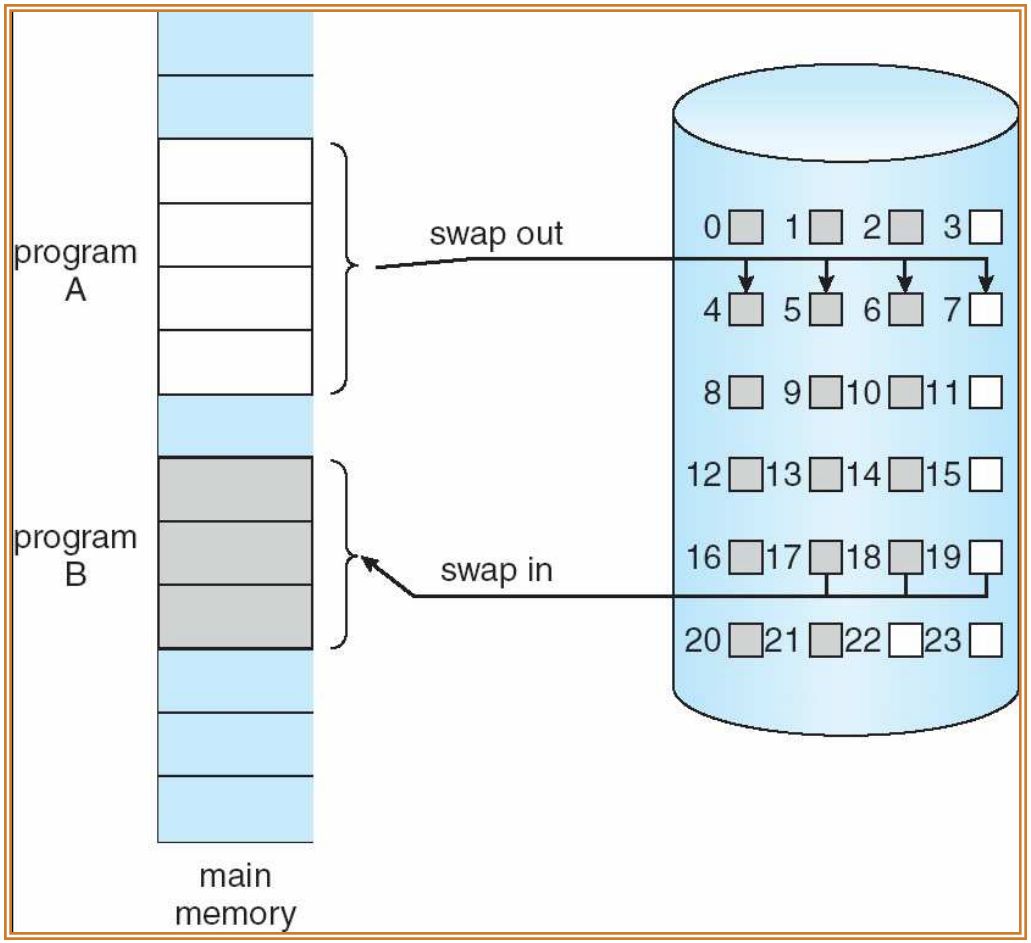

VIRTUAL MEMORY

Demand Paging

One instruction may require several pages. For example, a block move of data.

May page fault part way through an operation - may have to undo what was done. Example: an instruction crosses a page boundary.

Time to service page faults demands that they happen only infrequently.

Note here that the page table requires a "resident" bit showing that page is/isn't in memory. (Book uses "valid" bit to indicate residency. An "invalid" page is that way because a legal page isn't resident or because the address is illegal.

It makes more sense to have two bits - one indicating that the page is legal (valid) and a second to show that the page is in memory.

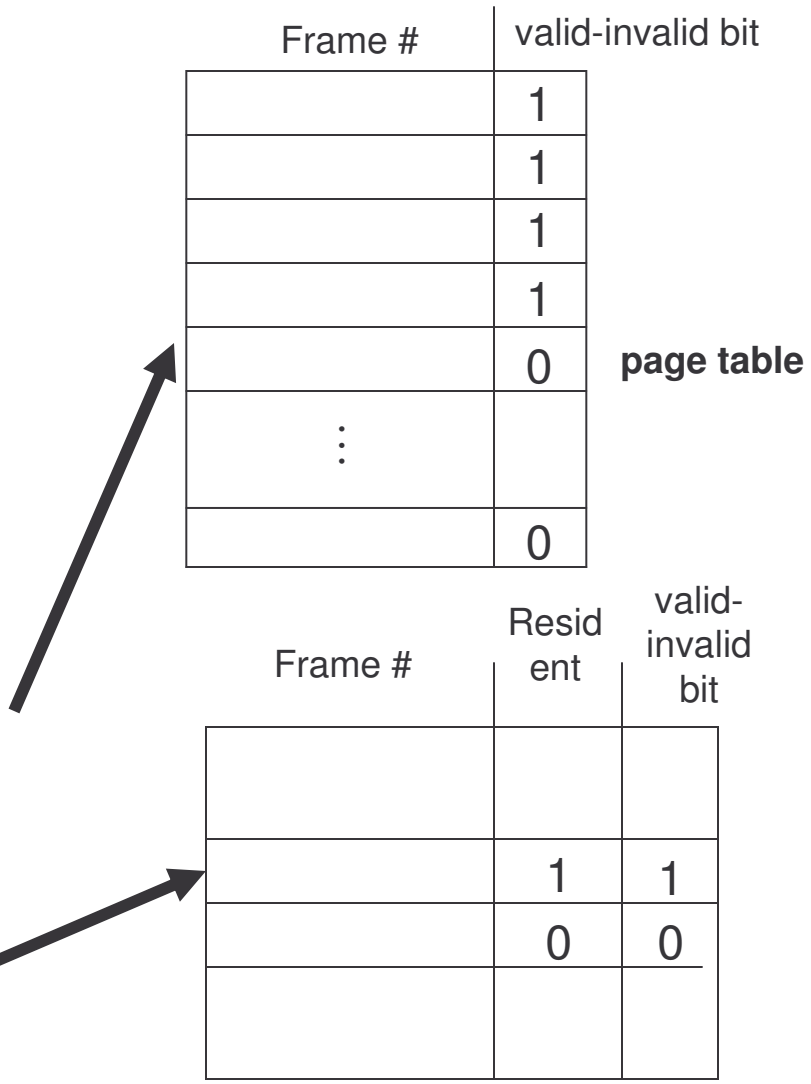

VIRTUAL MEMORY

Demand Paging

STEPS IN HANDLING A PAGE FAULT

1. The process has touched a page not currently in memory.
2. Check an internal table for the target process to determine if the reference was valid (do this in hardware.)
3. If page valid, but page not resident, try to get it from secondary storage.
4. Find a free frame; a page of physical memory not currently in use. (May need to free up a page.)
5. Schedule a disk operation to read the desired page into the newly allocated frame.
6. When memory is filled, modify the page table to show the page is now resident.
7. Restart the instruction that failed

Do these steps using the figure you can see on the next page.

VIRTUAL MEMORY

Demand Paging

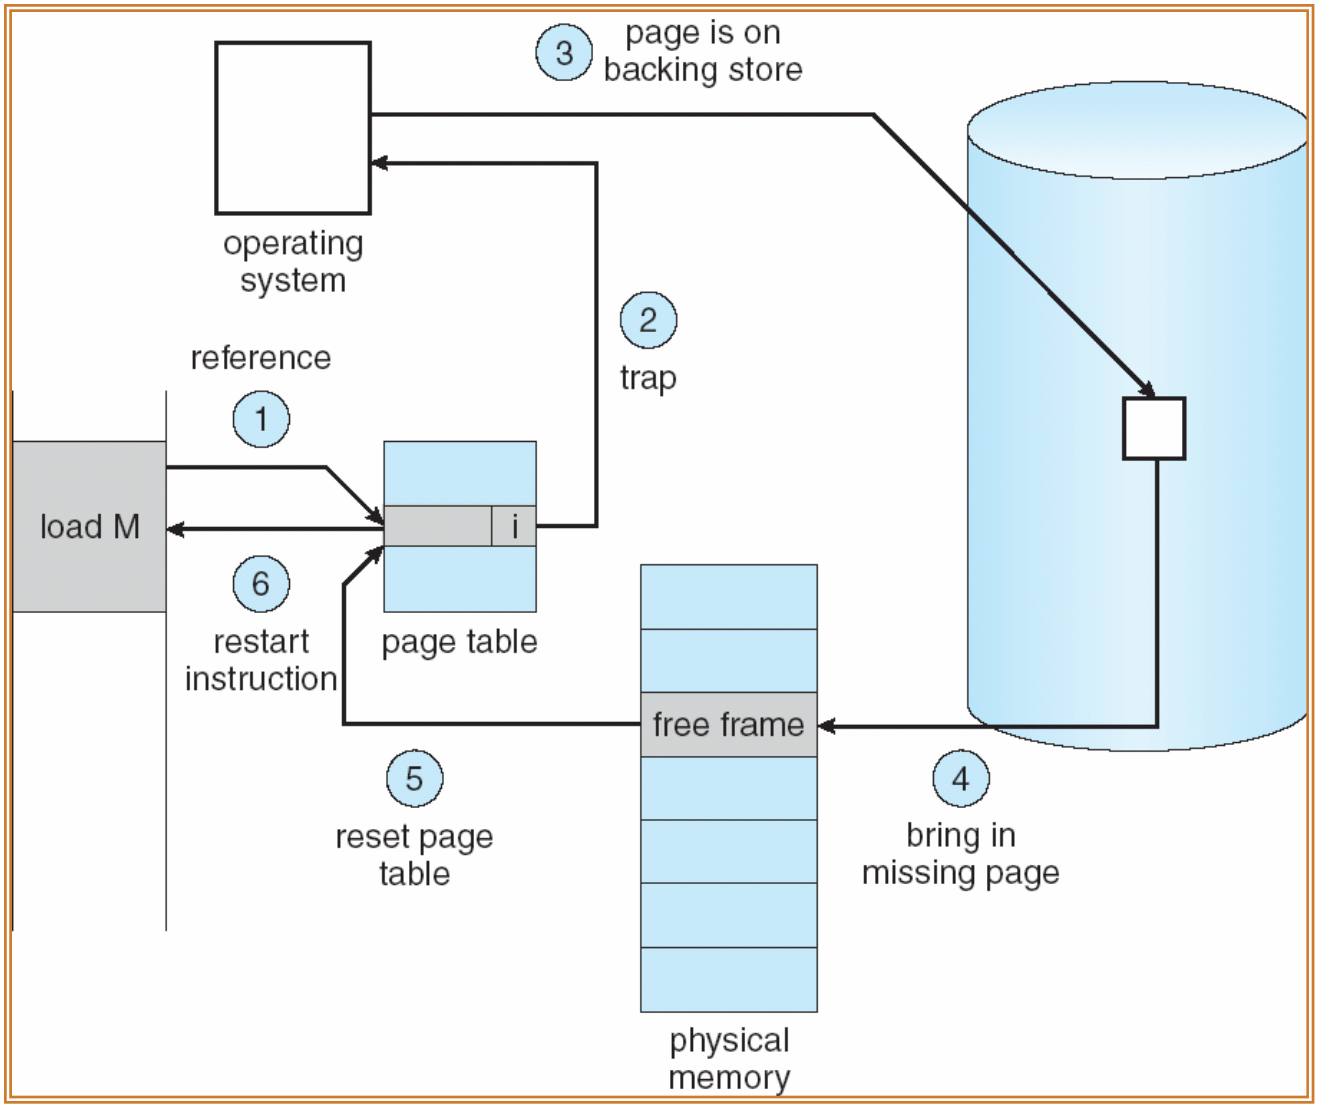

VIRTUAL MEMORY

Demand Paging

REQUIREMENTS FOR DEMAND PAGING (HARDWARE AND SOFTWARE) INCLUDE:

Page table mechanism

Secondary storage (disk or network mechanism.)

Software support for fault handlers and page tables.

Architectural rules concerning restarting of instructions. (For instance, block moves across faulted pages.)

VIRTUAL MEMORY

Demand Paging

PERFORMANCE OF DEMAND PAGING

We are interested in the effective access time: a combination of "normal" and "paged" accesses.

It's important to keep fraction of faults to a minimum. If fault ratio is "p", then

effective_access_time = (1 - p) * memory_access_time
+ p * page_fault_time.

Calculate the time to do a fault as shown in the text:

fault time = 10 milliseconds (why)
normal access = 100 nanoseconds (why)

How do these fit in the formula?

VIRTUAL MEMORY

The Picture When All Pages Are Not In Memory

Some of the pages belonging to this process are in memory, and some are on the disk.

A bit in the page table tells where to find the page.

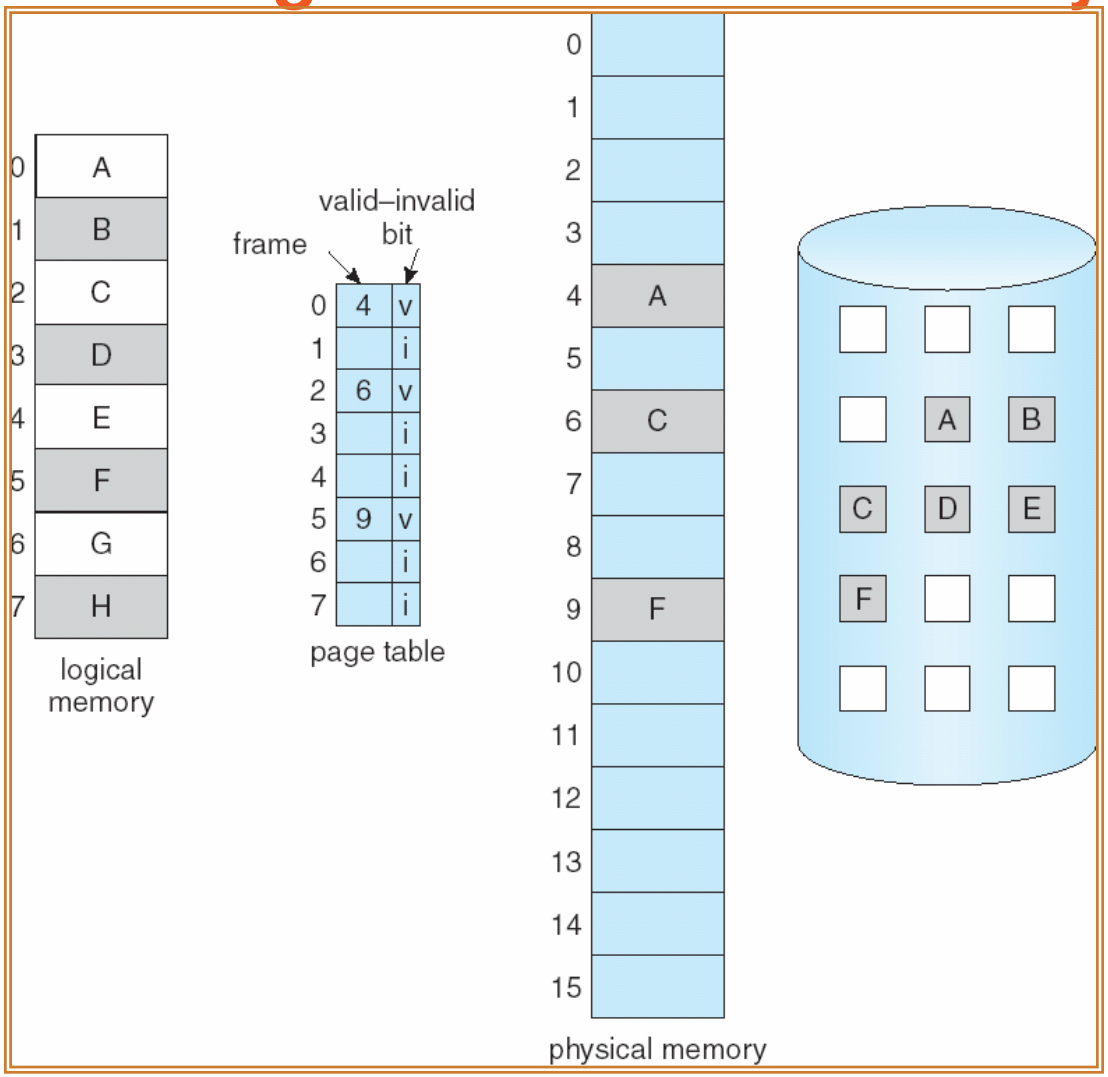

VIRTUAL MEMORY

Page Replacement

When we over-allocate memory, we need to push out something already in memory. Over-allocation may occur when programs need to fault in more pages than there are physical frames to handle.

Approach: If no physical frame is free, find one not currently being touched and free it. Steps to follow are:

1. Find requested page on disk.
2. Find a free frame.
   a. If there's a free frame, use it
   b. Otherwise, select a victim page.
   c. Write the victim page to disk.
3. Read the new page into freed frame. Change page and frame tables.
4. Restart user process.

Hardware requirements include "dirty" or modified bit.

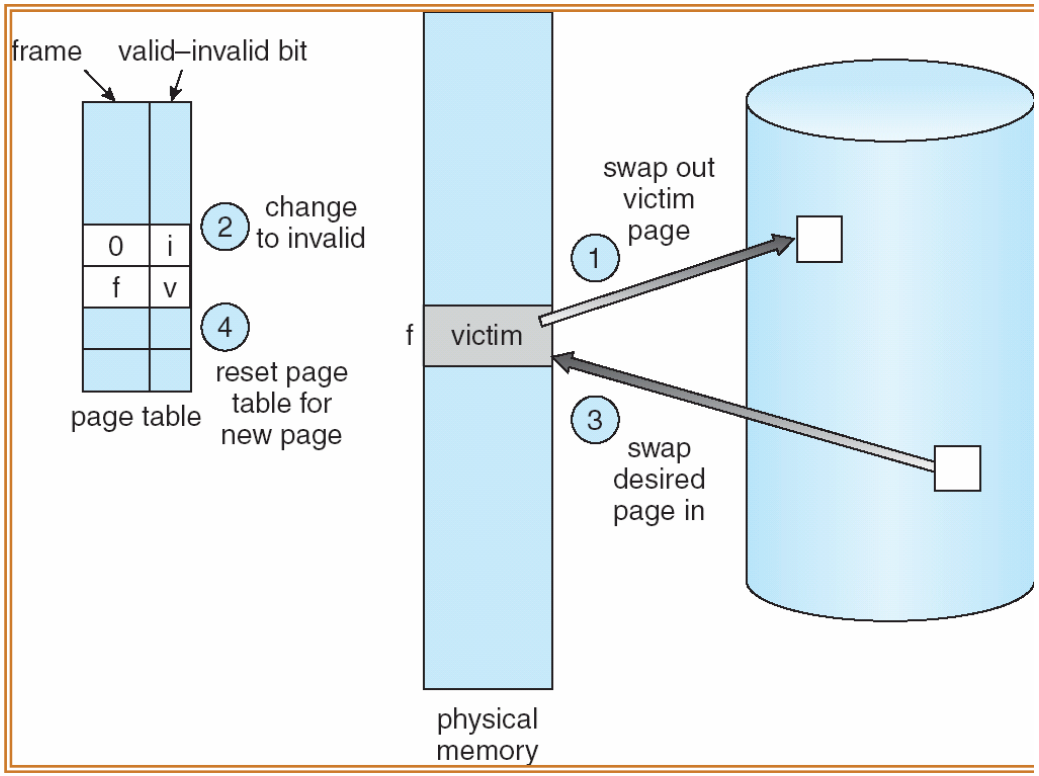

VIRTUAL MEMORY

Page Replacement

PAGE REPLACEMENT ALGORITHMS:

When memory is overallocated, we can either swap out some process, or overwrite some pages. Which pages should we replace?? <--- here the goal is to minimize the number of faults.

Here is an example reference string we will use to evaluate fault mechanisms:

Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

FIFO

Conceptually easy to implement; either use a time-stamp on pages, or organize on a queue. (The queue is by far the easier of the two methods.)

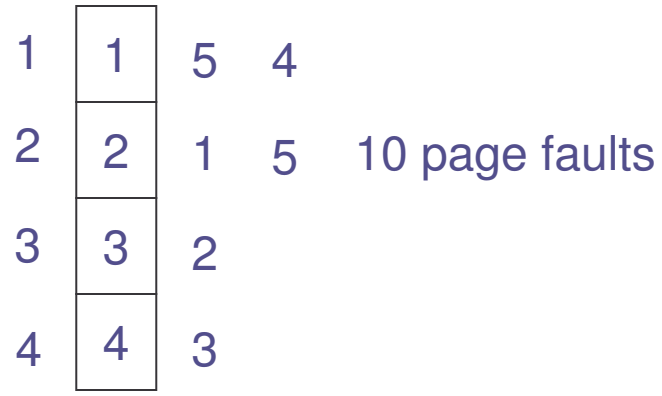

OPTIMAL REPLACEMENT

- This is the replacement policy that results in the lowest page fault rate.
- Algorithm: Replace that page which will not be next used for the longest period of time.
- Impossible to achieve in practice; requires crystal ball.

Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

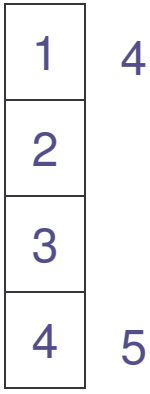

6 page faults

VIRTUAL MEMORY

Page Replacement

LEAST RECENTLY USED (LRU)

- Replace that page which has not been used for the longest period of time.
- Results of this method considered favorable. The difficulty comes in making it work.
- Implementation possibilities:

  Time stamp on pages - records when the page is last touched.
  Page stack - pull out touched page and put on top

Both methods need hardware assist since the update must be done on every instruction. So in practice this is rarely done.

Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

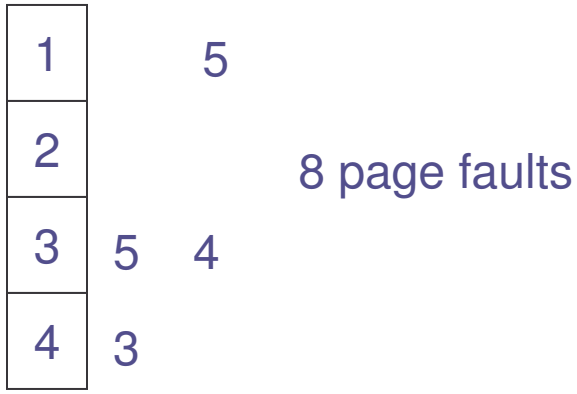

VIRTUAL MEMORY

Page Replacement

PAGE REPLACEMENT ALGORITHMS :

Using another string:

FIFO

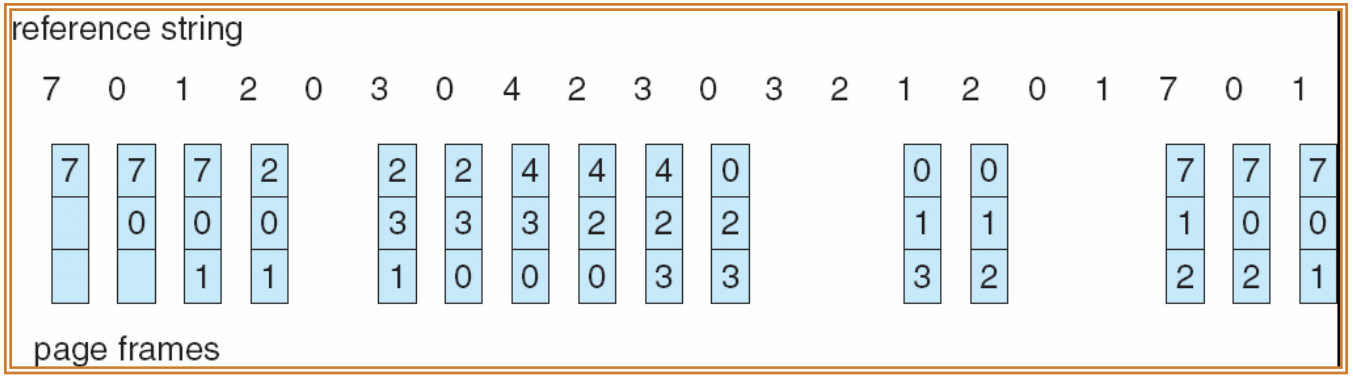

OPTIMAL

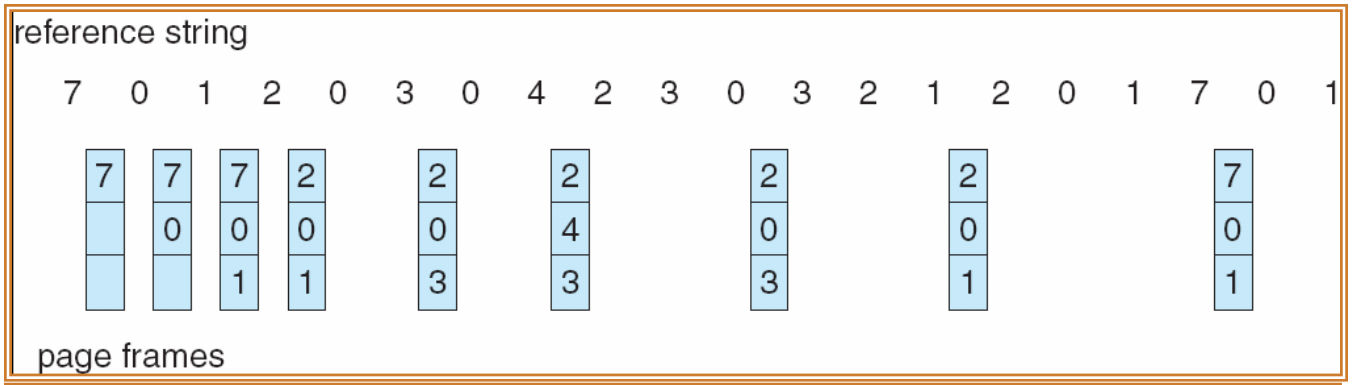

LRU

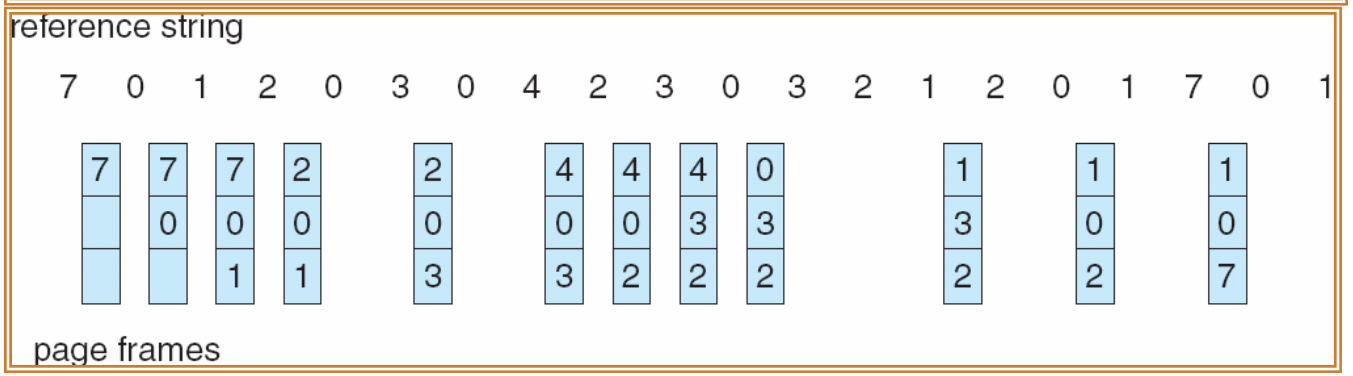

VIRTUAL MEMORY

Page Replacement

LRU APPROXIMATION

Uses a reference bit set by hardware when the page is touched. Then when a fault occurs, pick a page that hasn't been referenced.

Additional reference bits can be used to give some time granularity. Then pick the page with the oldest timestamp.

Second chance replacement: pick a page based on FIFO. If its reference bit is set, give it another chance. Envision this as a clock hand going around a circular queue. The faster pages are replaced, the faster the hand goes.

Maintain a modified bit, and preferentially replace unmodified pages.

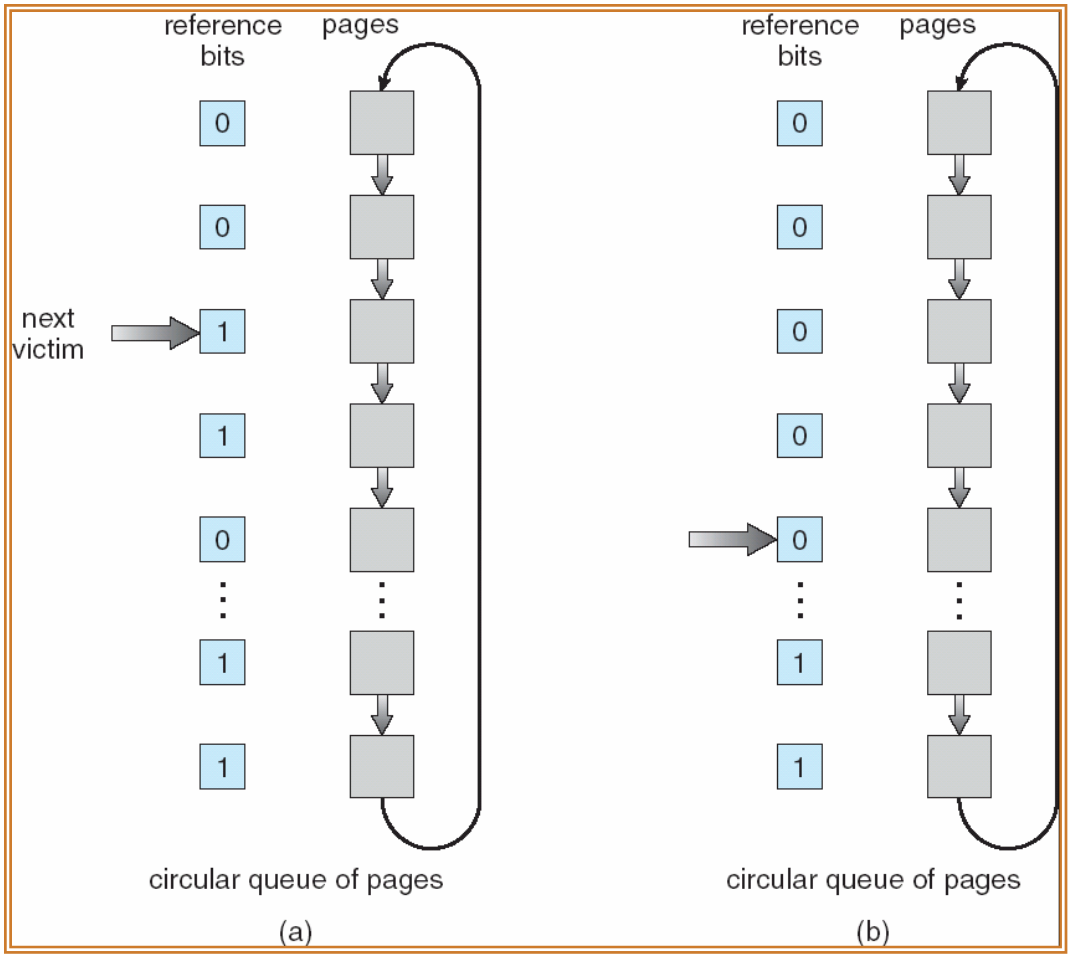

Second-Chance (clock) Page-Replacement Algorithm

VIRTUAL MEMORY

Page Replacement

ADD HOC (OR ADD-ON) ALGORITHMS

These methods are frequently used over and above the standard methods given above.

Maintain pools of free frames; write out loser at leisure

Occasional writes of dirties - make clean and then we can use them quickly when needed.

Write out and free a page but remember a page id in case the page is needed again - even though a page is in the free pool, it can be recaptured by a process (a soft page fault.)

VIRTUAL MEMORY

Page Allocation

ALLOCATION OF FRAMES:

What happens when several processes contend for memory? What algorithm determines which process gets memory - is page management a global or local decision?

A good rule is to ensure that a process has at least a minimum number of pages. This minimum ensures it can go about its business without constantly thrashing.

ALLOCATION ALGORITHMS

Local replacement -- the process needing a new page can only steal from itself. (Doesn't take advantage of entire picture.)

Global replacement - sees the whole picture, but a memory hog steals from everyone else

Can divide memory equally, or can give more to a needier process. Should high priority processes get more memory?

VIRTUAL MEMORY

Thrashing

Suppose there are too few physical pages (less than the logical pages being actively used). This reduces CPU utilization, and may cause increase in multiprogramming needs defined by locality.

A program will thrash if all pages of its locality aren't present in the working set.

Two programs thrash if they fight each other too violently for memory.

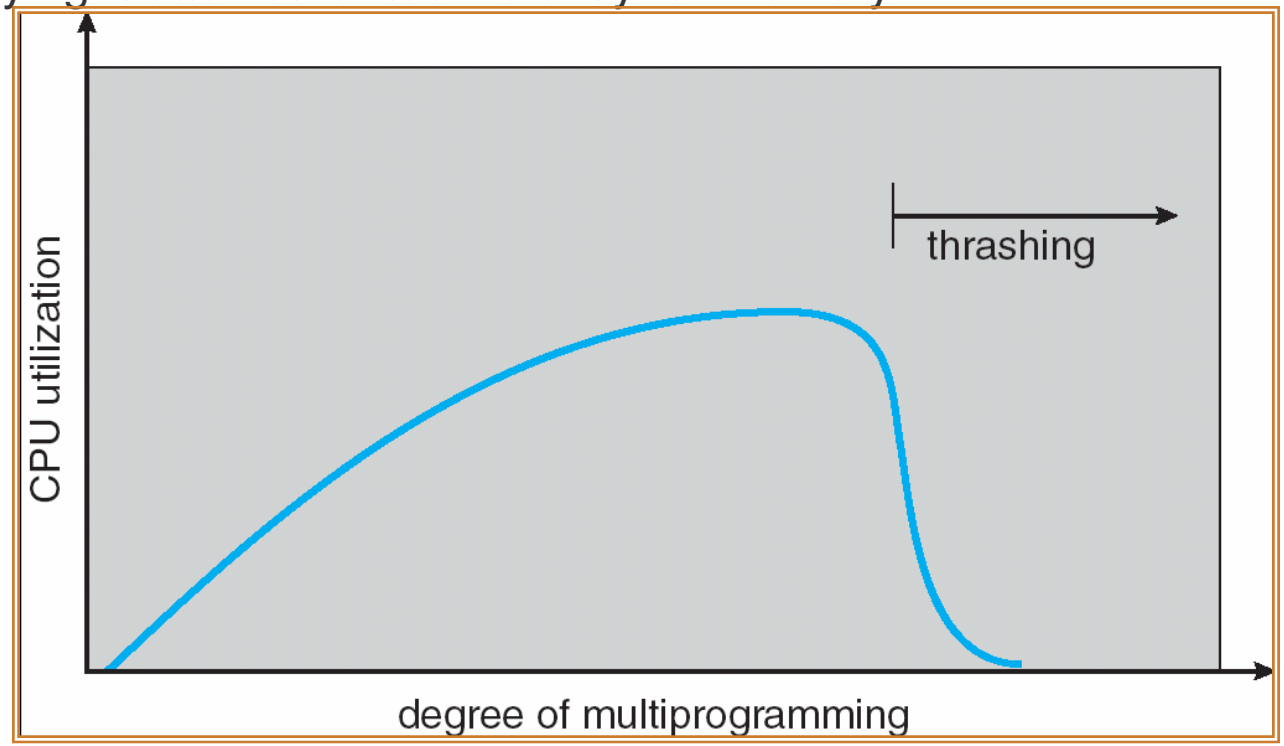

VIRTUAL MEMORY

Locality of Reference

Locality of reference: Programs access memory near where they last accessed it.

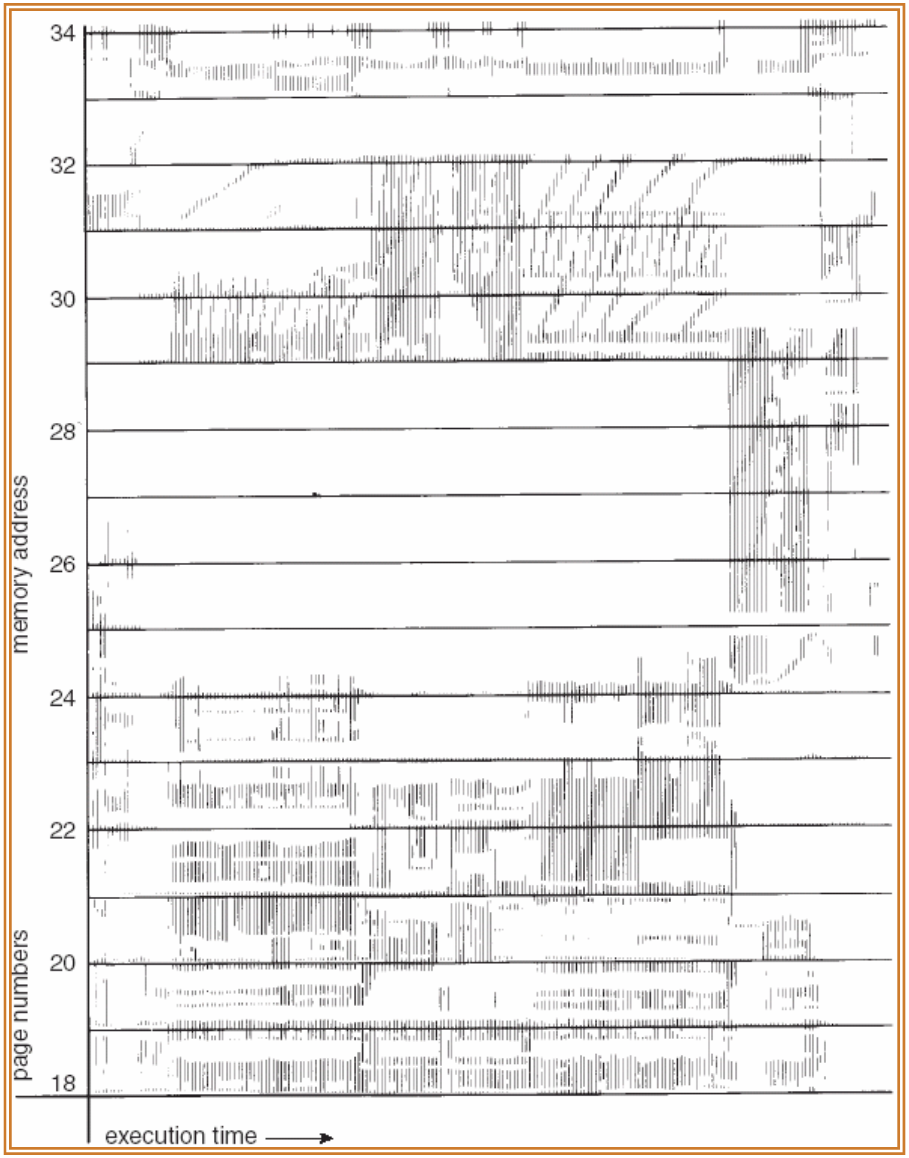

WORKING SET MODEL

The pages used by a process within a window of time are called its working set.

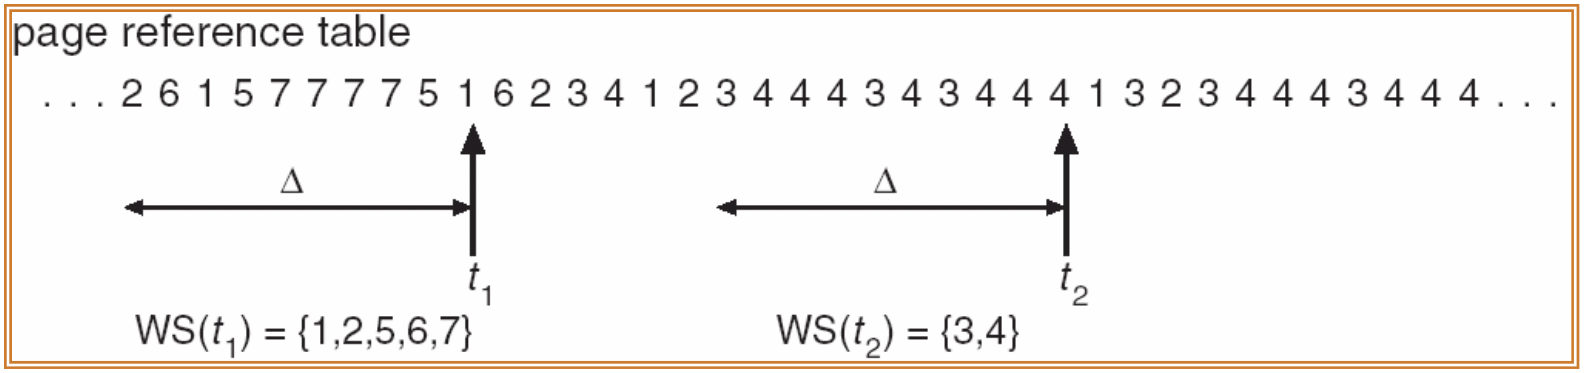

Changes continuously - hard to maintain an accurate number. How can the system use this number to give optimum memory to the process?

VIRTUAL MEMORY

Working Set Model

PAGE FAULT FREQUENCY

This is a good indicator of thrashing. If the process is faulting heavily, allocate it more frames. If faulting very little, take away some frames.

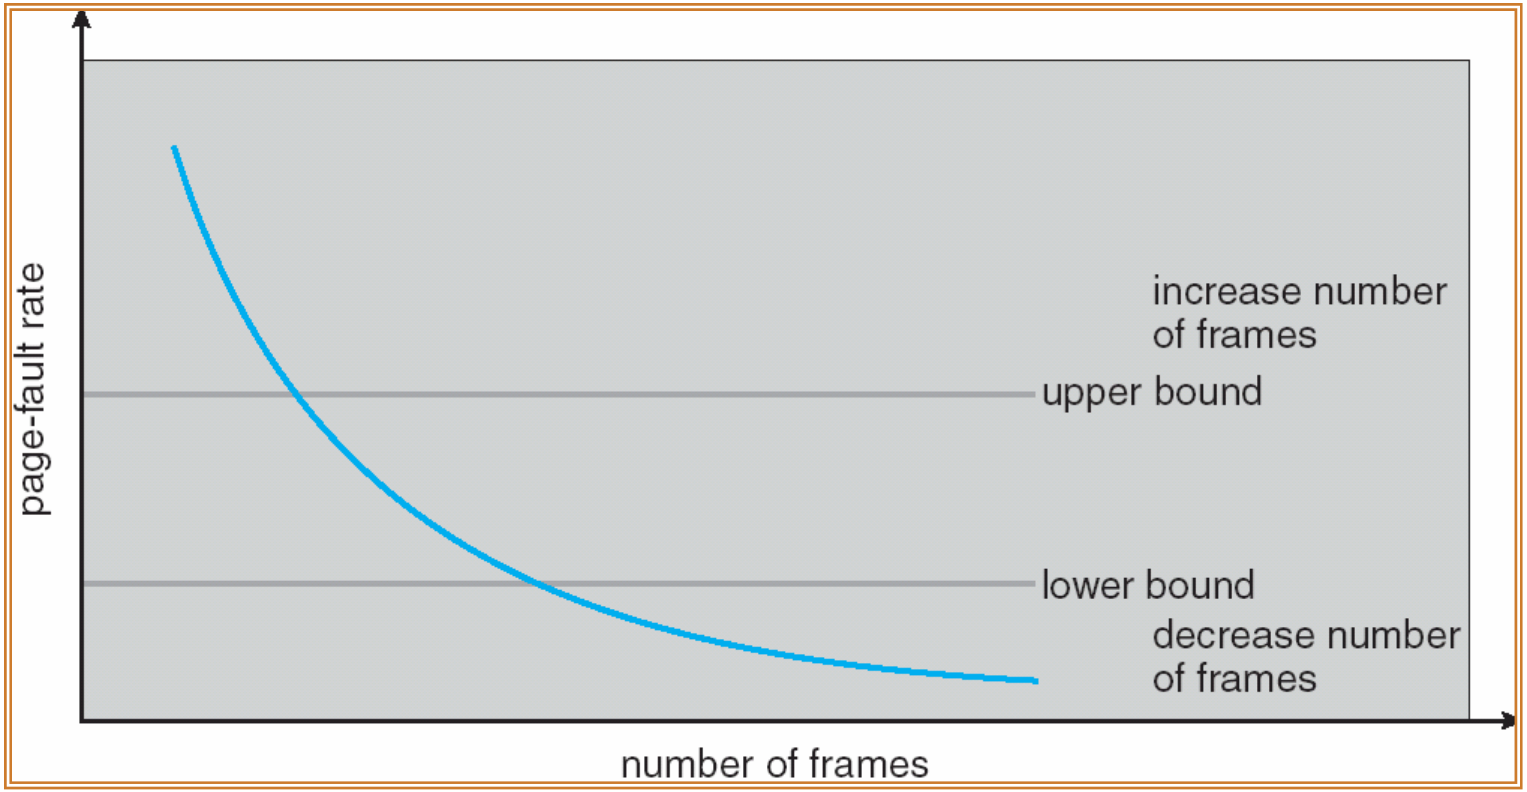

VIRTUAL MEMORY Other Issues

PREPAGING

- Bring lots of pages into memory at one time, either when the program is initializing, or when a fault occurs.
- Uses the principle that a program often uses the page right after the one previously accessed (locality of reference.)

PAGE SIZE

- If too big, there's considerable fragmentation and unused portions of the page.
- If too small, table maintenance is high and so is I/O.
- Has ramifications in code optimization.

VIRTUAL MEMORY

Other Issues

Memory Mapped IO

- Allows file I/O to be treated as routine memory access by **mapping** a disk block to a page in memory
- A file is initially read using demand paging. Multiple page-sized portions of the file are read from the file system into physical pages. Subsequent reads/writes to/from the file are treated as ordinary memory accesses.
- Simplifies file access by treating file I/O through memory rather than **read()** **write()** system calls
- Also allows several processes to map the same file allowing the pages in memory to be shared

VIRTUAL MEMORY

Other Issues

Memory Mapped IO

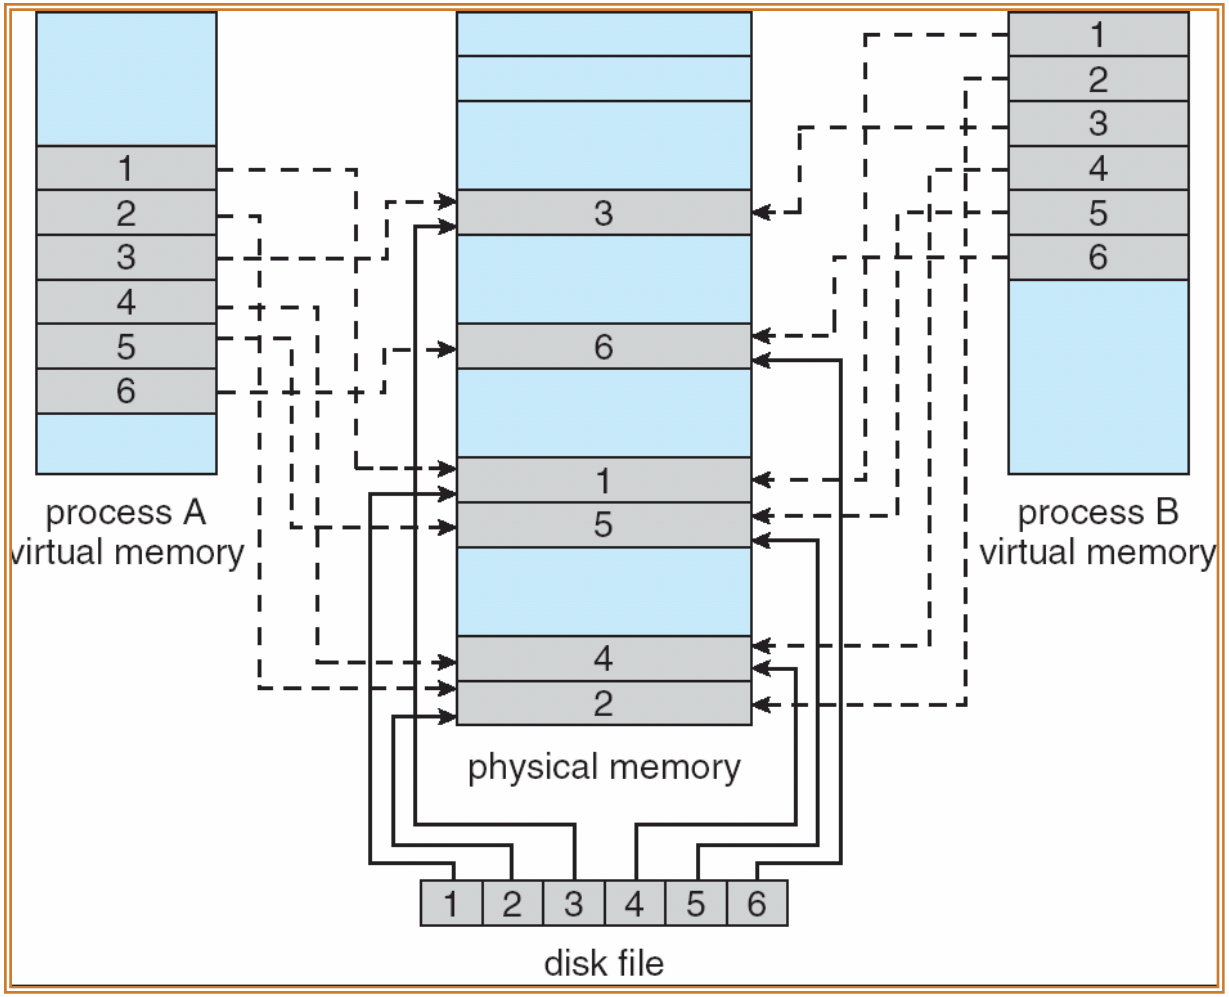

VIRTUAL MEMORY

Other Issues

Program structure

- int data [128,128];
- Each row is stored in one page
- Program 1

```
for (j = 0; j <128; j++)
    for (i = 0; i < 128; i++)
        data[i,j] = 0;
```

128 x 128 = 16,384 page faults

- Program 2

```
for (i = 0; i < 128; i++)
    for (j = 0; j < 128; j++)
        data[i,j] = 0;
```

128 page faults

Which method should you program??

VIRTUAL MEMORY

Wrap up

Virtual memory is how we stuff large programs into small physical memories.

We perform this magic by using demand paging, to bring in pages only when they are needed.

But to bring pages into memory, means kicking other pages out, so we need to worry about paging algorithms.