

COP 4610

Operating System Principles

Virtual Memory

1

Memory Pages

- Why do we have pages?
- What are advantages & disadvantages of using “page model”?

2

Background

- Is all of your code used?
 - Error code
 - Unusual routines, certain options/features
 - Large data structures, lists, arrays, ...
- Entire program code not needed at same time!

What if we could only partially load what we need?

3

Virtual Memory


Separation of user logical memory
from physical memory

- Can be implemented via **demand paging** & **demand segmentation**

4


Benefits

- **Huge accessible space**
 - Map virtual to physical
- Can easily **share** things
- **Efficient process creation**
 - Only load partial process
- **More programs** running
- **Less I/O** needed to load or swap processes
 - Some parts of process will never be needed



5

Virtual Memory >> Physical Memory



6

Virtual Address Space


- Enables **sparse** address spaces
- **System libraries** shared via mapping into virtual address space
- **Shared memory** by mapping pages read-write into virtual address space
- Pages can be shared during `fork()`, speeding **process creation**

COP 4610 – Operating System Principles

7

7

Shared Library Using Virtual Memory



COP 4610 – Operating System Principles

8

8

Demand Paging


- Load page **ONLY WHEN NEEDED**:
 - Less I/O needed
 - Less memory needed
 - Faster response
 - More users
- **Lazy swapper** – never swaps a page into memory unless page will be needed
 - Swapper that deals with pages is a **pager**

COP 4610 – Operating System Principles

9

9

Demand Paging




COP 4610 – Operating System Principles

10

10

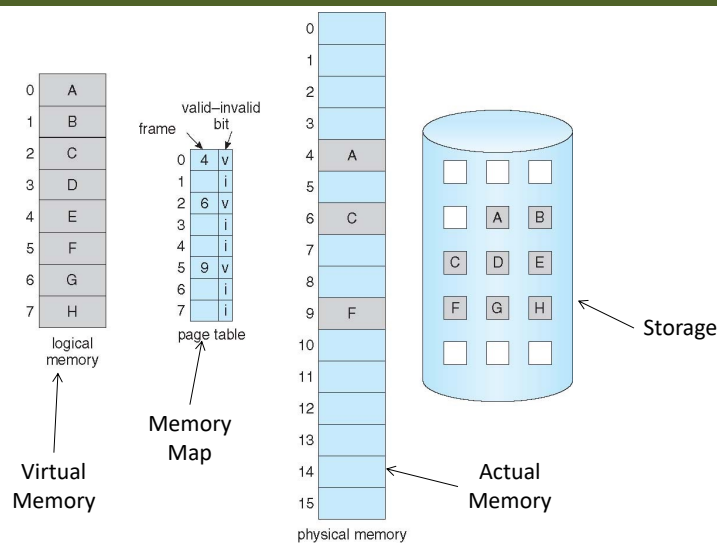
Valid-Invalid Bit

- With each page table entry, a valid–invalid bit is associated
 - **v** ⇒ in-memory (memory resident)
 - **i** ⇒ not-in-memory
- Initial valid–invalid bit
 - Set to **i** on all entries
- During **address translation**, if valid–invalid bit in page table entry is **i** ⇒ **page fault**



11

Page Table When Some Pages Are Not in Main Memory



12

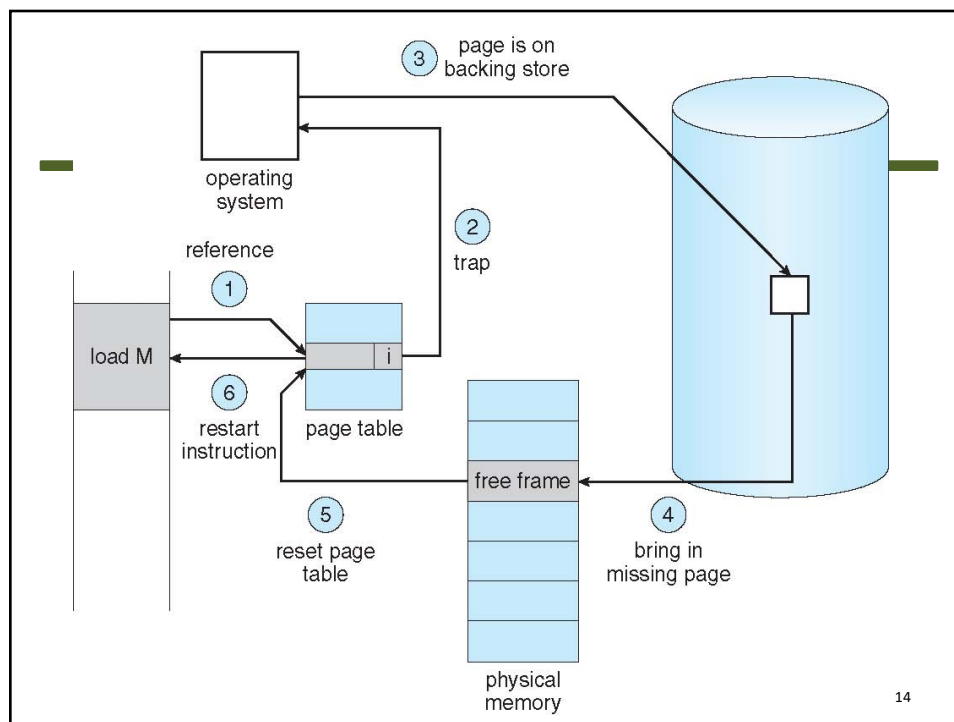
Page Fault

- If there is a reference to a page, first reference to that page will trap to operating system (page fault)
1. Operating system checks if it was an invalid reference (if so, abort)
 2. Get empty frame
 3. Swap page into frame
 4. Reset tables to indicate page now in memory
Set validation bit = **v**
 5. Restart the instruction that caused the page fault

COP 4610 – Operating System Principles

13

13



14

14

Stages in Demand Paging

1. Trap to the operating system
2. Save the user registers and process state
3. Determine that the interrupt was a page fault
4. Check that the page reference was legal and determine the location of the page on the disk
5. Issue a read from the disk to a free frame:
 1. Wait in a queue for this device until the read request is serviced
 2. Wait for the device seek and/or latency time
 3. Begin the transfer of the page to a free frame
6. While waiting, allocate the CPU to some other user
7. Receive an interrupt from the disk I/O subsystem (I/O completed)
8. Save the registers and process state for the other user
9. Determine that the interrupt was from the disk
10. Correct the page table and other tables to show page is now in memory
11. Wait for the CPU to be allocated to this process again
12. Restore the user registers, process state, and new page table, and then resume the interrupted instruction

COP 4610 – Operating System Principles

15

15

Performance of Demand Paging

- Page Fault Rate $0 \leq p \leq 1$
 - if $p = 0$ no page faults
 - if $p = 1$, every reference is a fault
- Effective Access Time (EAT)

$$\text{EAT} = (1-p) \cdot m + p \cdot \text{PFST}$$
 - m = memory access time
 - PFST = Page Fault Service Time (swapping in/out, restarting, etc.)

COP 4610 – Operating System Principles

16

16

Demand Paging Example

- Memory access time = 200 nanoseconds
- Average page-fault service time = 8 milliseconds
- $EAT = (1 - p) \times 200 + p (8 \text{ milliseconds})$
 $= (1 - p) \times 200 + p \times 8,000,000$
 $= 200 + p \times 7,999,800$
- If one access out of 1,000 causes a page fault, then
 $EAT = 8.2 \text{ microseconds}$.
 This is a slowdown by a factor of **40!!**
- If want performance degradation < 10 percent
 - $220 > 200 + 7,999,800 \times p$
 - $20 > 7,999,800 \times p$
 - $p < .0000025$
 - < one page fault in every 400,000 memory accesses


17

Process Creation: Copy-on-Write

- **Copy-on-Write (COW)** allows both parent and child processes to initially *share* the same pages in memory
 - If either process modifies a shared page, only then is the page copied
- COW allows more efficient process creation as only modified pages are copied
- In general, free pages are allocated from a **pool of zero-fill-on-demand** pages
 - Why zero-out a page before allocating it?
- `vfork()` variation on `fork()` system call has parent suspend and child using copy-on-write address space of parent
 - Designed to have child call `exec()`
 - Very efficient


18

Before Process 1 Modifies Page C



19

After Process 1 Modifies Page C



20

What Happens if There is no Free Frame?


- **Page replacement** – find some page in memory, but not really in use, page it out
 - Algorithm – terminate? swap out? replace the page?
 - Performance – want an algorithm which will result in minimum number of page faults
- Same page may be brought into memory several times

COP 4610 – Operating System Principles

21

21

Page Replacement



22

22

Page Replacement

- Prevent over-allocation of memory by modifying page-fault service routine to include page replacement
- Use **modify (dirty) bit** to reduce overhead of page transfers – only modified pages are written to disk

23

Basic Page Replacement

1. Find the location of the desired page on disk
2. Find a free frame:
 - If there is a free frame, use it
 - If there is no free frame, use page replacement algorithm to select the **victim frame**
 - Write victim frame to disk if dirty
3. Bring the desired page into the (newly) free frame; update the page and frame tables
4. Continue the process by restarting the instruction that caused the trap


24

Page Replacement Dilemma

- Raw Performance
 - If it is dirty, need to write it to disk (swap space)
 - Need to write the whole page
- Future Page Faults
 - The one swapped out could be needed soon again

25

Page Faults vs Number of Frames



26

Page and Frame Replacement Algorithms

- **Frame-allocation algorithm** determines
 - How many frames to give each process
 - Which frames to replace
- **Page-replacement algorithm**
 - Want lowest page-fault rate on both first access and re-access
- Evaluate algorithm by running it on a particular string of memory references (reference string) and computing the number of page faults on that string
 - String is just page numbers, not full addresses
 - Repeated access to the same page does not cause a page fault
- In all our examples, the reference string is


7,0,1,2,0,3,0,4,2,3,0,3,0,3,2,1,2,0,1,7,0,1

27

FIFO Page Replacement

3 pages available **7,0,1,2,0,3,0,4,2,3,0,3,0,3,2,1,2,0,1,7,0,1**

reference string



page frames


Do more available pages always improve performance?

Possible to make things worse – **Belady's Anomaly**

28

Belady's Anomaly

- 1,2,3,4,1,2,5,1,2,3,4,5
- 3 frames (9 page faults) versus 4 frames (10 page faults)



COP 4610 – Operating System Principles

29

29

Optimal Algorithm (OPT)

- Replace page that will not be used for longest period of time
- How do you know this?
- Used for measuring how well your algorithm performs

COP 4610 – Operating System Principles

30

30

Optimal Page Replacement

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	2	2	2	7
	0	0	0	0	4	0	0	0
	1	1	1	3	3	3	1	1

page frames

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

FIFO

7	7	7	2	2	2	4	4	4	0	0	0	7	7	7
	0	0	0	3	3	3	2	2	2	1	1	1	0	0
	1	1	1	1	0	0	0	3	3	3	2	2	2	1

page frames

COP 4610 – Operating System Principles

31

31

Least Recently Used (LRU) Algorithm

- Use past knowledge rather than future
- Replace page that has not been used in the most amount of time
- Associate time of last use with each page

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	4	4	4	0	1	1	1
	0	0	0	0	0	0	3	3	3	0	0
	1	1	1	3	3	2	2	2	2	2	7

page frames

- 12 faults – better than FIFO but worse than OPT
- Generally good algorithm and frequently used
- But how to implement?

COP 4610 – Operating System Principles

32

32

LRU Algorithm


- Stack implementation
 - Keep a stack of page numbers in a double link form
 - Page referenced:
 - move it to the top
 - requires 6 pointers to be changed
 - But each update more expensive
 - No search for replacement
- LRU and OPT are cases of **stack algorithms** that don't have Belady's Anomaly

COP 4610 – Operating System Principles

33

33

Stack Approach



COP 4610 – Operating System Principles

34

34

LRU Approximation Algorithms


- **Reference bit**
 - With each page associate a bit, initially = 0
 - When page is referenced bit set to 1
 - Replace any with reference bit = 0 (if one exists)
 - We do not know the order, however
- **Second-chance algorithm**
 - Generally FIFO, plus hardware-provided reference bit
 - **Clock** replacement
 - If page to be replaced has
 - Reference bit = 0 -> replace it
 - Reference bit = 1 then:
 - set reference bit 0, leave page in memory
 - replace next page, subject to same rules

COP 4610 – Operating System Principles

35

35

LRU Approximation Algorithms



COP 4610 – Operating System Principles

36

36

Allocation of Frames

- Each process needs *minimum* number of frames
- Example: IBM 370 – 6 pages to handle SS MOVE instruction:
 - instruction is 6 bytes, might span 2 pages
 - 2 pages to handle *from*
 - 2 pages to handle *to*
- *Maximum* of course is total frames in the system
- Two major allocation schemes
 - fixed allocation
 - priority allocation
- Many variations

COP 4610 – Operating System Principles

37

37

Fixed Allocation

- **Equal allocation** – For example, if there are 100 frames (after allocating frames for the OS) and 5 processes, give each process 20 frames
 - Keep some as free frame buffer pool
- **Proportional allocation** – Allocate according to the size of process
 - Dynamic as degree of multiprogramming, process sizes change

COP 4610 – Operating System Principles

38

38

Fixed Allocation

s_i □ size of process p_i

S □ $\sum s_i$

m □ total number of frames

a_i □ allocation for p_i □ $\frac{s_i}{S} \times m$

m □ 64

s_1 □ 10

s_2 □ 127

a_1 □ $\frac{10}{137} \times 64 \approx 5$

a_2 □ $\frac{127}{137} \times 64 \approx 59$

39

Priority Allocation

- Use a proportional allocation scheme using priorities rather than size
- If process P_i generates a page fault,
 - select for replacement one of its frames
 - select for replacement a frame from a process with lower priority number

40

Global vs. Local Allocation

- **Global replacement** – process selects a replacement frame from the set of all frames; one process can take a frame from another
 - But then process execution time can vary greatly
 - But greater throughput, so more common


- **Local replacement** – each process selects from only its own set of allocated frames
 - More consistent per-process performance
 - But possibly underutilized memory

Thrashing

- If a process does not have “enough” pages, the page-fault rate is very high
 - Page fault to get page
 - Replace existing frame
 - But quickly need replaced frame back
 - This leads to:
 - Low CPU utilization
 - Operating system thinking that it needs to increase the degree of multiprogramming
 - Another process added to the system

- **Thrashing** \equiv a process is busy swapping pages in and out

Thrashing (Cont.)



COP 4610 – Operating System Principles

43

43

Demand Paging and Thrashing


- Why does demand paging work?
 - Locality model**
 - Process migrates from one locality to another
 - Localities may overlap
- Why does thrashing occur?
 - Σ size of locality > total memory size
 - Limit effects by using local or priority page replacement

COP 4610 – Operating System Principles

44

44

Locality In A Memory-Reference Pattern




COP 4610 – Operating System Principles

45

45

Page-Fault Frequency

- If actual rate too low, process may lose frames
- If actual rate too high, process may gain frames



COP 4610 – Operating System Principles

46

46

Other Considerations -- Prepaging

- Prepaging
 - To reduce the large number of page faults that occurs at process startup
 - Prepage all or some of the pages a process will need, before they are referenced
 - But if prepagged pages are unused, I/O and memory was wasted
 - Assume s pages are prepagged and α of the pages is used
 - Is cost of $s * \alpha$ saved page faults > or < than the cost of prepagging $s * (1 - \alpha)$ unnecessary pages?
 - α near zero \Rightarrow prepagging loses

COP 4610 – Operating System Principles

47

47

Other Issues – Page Size

- Sometimes OS designers have a choice
 - Especially if running on custom-built CPU
- Page size selection must take into consideration:
 - Fragmentation
 - Page table size
 - I/O overhead
 - Number of page faults
 - Locality
 - TLB size and effectiveness
- Always power of 2, usually in the range 2^{12} (4,096 bytes) to 2^{22} (4,194,304 bytes)

COP 4610 – Operating System Principles

48

48

Other Issues – TLB Reach

- TLB Reach - The amount of memory accessible from the TLB
- $TLB\ Reach = (TLB\ Size) \times (Page\ Size)$
- Ideally, the working set of each process is stored in the TLB
 - Otherwise there is a high degree of page faults
- Increase the Page Size
 - This may lead to an increase in fragmentation as not all applications require a large page size
- Provide Multiple Page Sizes
 - This allows applications that require larger page sizes the opportunity to use them without an increase in fragmentation

49

Other Issues – Program Structure

- Program structure
 - `int[128,128] data;`
 - 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

50