


Thrashing

- ▶ If a process does not have “enough” pages, the page-fault rate is very high. This leads to:
 - low CPU utilization
 - operating system thinks that it needs to increase the degree of multiprogramming because of low cpu utilization
 - another process added to the system
- ▶ Thrashing $\equiv$ a process is busy swapping pages in and out



Thrashing (Cont.)



Demand Paging and Thrashing

- ▶ Why does demand paging work?

Locality model

- Process migrates from one locality to another
- Localities may overlap
- E.g.

```
for (.....) {  
    computations;  
}  
  
.....  
for (..... ) {  
    computations;  
}
```

- ▶ Why does thrashing occur?
 Σ size of locality > total memory size



Working-Set Model


- ▶ $\Delta \equiv$ working-set window $\equiv$ a fixed number of page references
Example: 10,000 instruction
- ▶ WSS_i (working set of Process P_i) = total number of pages referenced in the most recent Δ (varies in time)
 - if Δ too small will not encompass entire locality
 - if Δ too large will encompass several localities
 - if $\Delta = \infty \Rightarrow$ will encompass entire program
- ▶ $D = \sum WSS_i \equiv$ total demand frames
- ▶ if $D > m \Rightarrow$ Thrashing
- ▶ Policy if $D > m$, then suspend one of the processes




Working-set model

page reference table

... 2 6 1 5 7 7 7 7 5 1 6 2 3 4 1 2 3 4 4 4 3 4 3 4 4 4 1 3 2 3 4 4 4 3 4 4 4 ...



$WS(t_1) = \{1, 2, 5, 6, 7\}$



$WS(t_2) = \{3, 4\}$




Keeping Track of the Working Set

- ▶ Approximate with interval timer + a reference bit
- ▶ Example: $\Delta = 10,000$
 - Timer interrupts after every 5000 time units
 - Keep in memory 2 bits for each page
 - Whenever a timer interrupts copy and sets the values of all reference bits to 0
 - If one of the bits in memory = 1 $\Rightarrow$ page in working set
- ▶ Why is this not completely accurate?
- ▶ Improvement = 10 bits and interrupt every 1000 time units



Page-Fault Frequency Scheme

- ▶ Establish “acceptable” page-fault rate
 - If actual rate too low, process loses frame
 - If actual rate too high, process gains frame



Other Issues -- 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 prepaged pages are unused, I/O and memory was wasted
- Assume s pages are prepaged and α of the pages is used
 - Is cost of $s * \alpha$ save pages faults $>$ or $<$ than the cost of prepaging
 $s * (1 - \alpha)$ unnecessary pages?
 - α near zero $\Rightarrow$ prepaging loses



Other Issues – Page Size

- ▶ Page size selection must take into consideration:
 - fragmentation
 - table size
 - I/O overhead
 - locality



Other Issues – TLB Reach

- ▶ TLB Reach - The amount of memory accessible from the TLB
- ▶ $\text{TLB Reach} = (\text{TLB Size}) \times (\text{Page Size})$
- ▶ 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.



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



Wrapup

- ▶ Memory hierarchy:
 - Speed: L1, L2, L3 caches, main memory, disk etc.
 - Cost: disk, main memory, L3, L2, L1 etc.
- ▶ achieve good speed by moving “interesting” objects to higher cache levels while moving “uninteresting” objects to lower cache levels
- ▶ Hardware provides reference bit, modify bit, page access counters, page table validity bits
- ▶ OS sets them appropriately such that it will be notified via page fault
 - OS provides policies
 - Hardware provides mechanisms
- ▶ Implement VM, COW etc. that are tuned to observed workloads



Chapter 10: File-System Interface

► Objectives:

- To explain the function of file systems
- To describe the interfaces to file systems
- To discuss file-system design tradeoffs, including access methods, file sharing, file locking, and directory structures
- To explore file-system protection



File Concept

- ▶ Contiguous persistent logical address space, can be storing data or programs
- ▶ File Structure:
 - None - sequence of words, bytes
 - Simple record structure
 - Lines
 - Fixed length
 - Variable length
 - Complex Structures
 - Formatted document
 - Relocatable load file
- ▶ Can simulate last two with first method by inserting appropriate control characters
- ▶ Who decides:
 - Operating system
 - Program



File Types – Name, Extension

file type	usual extension	function
executable	exe, com, bin or none	ready-to-run machine-language program
object	obj, o	compiled, machine language, not linked
source code	c, cc, java, pas, asm, a	source code in various languages
batch	bat, sh	commands to the command interpreter
text	txt, doc	textual data, documents
word processor	wp, tex, rtf, doc	various word-processor formats
library	lib, a, so, dll	libraries of routines for programmers
print or view	ps, pdf, jpg	ASCII or binary file in a format for printing or viewing
archive	arc, zip, tar	related files grouped into one file, sometimes compressed, for archiving or storage
multimedia	mpeg, mov, rm, mp3, avi	binary file containing audio or A/V information



File Attributes


- ▶ **Name** – only information kept in human-readable form
- ▶ **Identifier** – unique tag (number) identifies file within file system
- ▶ **Type** – needed for systems that support different types
- ▶ **Location** – pointer to file location on device
- ▶ **Size** – current file size
- ▶ **Protection** – controls who can do reading, writing, executing
- ▶ **Time, date, and user identification** – data for protection, security, and usage monitoring
- ▶ Information about files are kept in the directory structure, which is maintained on the disk



Examples

► UNIX: ls -li

```
26047823 -rw-r--r--  1 surendar staff
596480 Mar 16 20:17 Lecture22.ppt
```



File Operations

- ▶ File is an **abstract data type**
- ▶ **File operations:**
 - Create
 - Write
 - Read
 - Reposition within file (seek)
 - Delete
 - Truncate
- ▶ *Open(F_i)* – search the directory structure on disk for entry F_i , and move the content of entry to memory
- ▶ *Close (F_i)* – move the content of entry F_i in memory to directory structure on disk



Open Files

- ▶ Several pieces of data are needed to manage open files:
 - File pointer: pointer to last read/write location, per process that has the file open
 - File-open count: counter of number of times a file is open – to allow removal of data from open-file table when last processes closes it
 - Disk location of the file: cache of data access information
 - Access rights: per-process access mode information



Open File Locking

- ▶ Provided by some operating systems and file systems
- ▶ Mediates access to a file
- ▶ Mandatory or advisory:
 - **Mandatory** – access is denied depending on locks held and requested
 - **Advisory** – processes can find status of locks and decide what to do



Access Methods

► Sequential Access

read next
write next
reset
no read after last write
(rewrite)


► Direct Access

read n
write n
position to n
read next
write next
rewrite n

n = relative block number



Sequential-access File




Simulation of Sequential Access on a Direct-access File

sequential access	implementation for direct access
<i>reset</i>	<i>cp = 0;</i>
<i>read next</i>	<i>read cp;</i> <i>cp = cp + 1;</i>
<i>write next</i>	<i>write cp;</i> <i>cp = cp + 1;</i>



Directory Structure


- ▶ A collection of nodes containing information about all files



Both the directory structure and the files reside on disk
Backups of these two structures are kept on tapes



A Typical File-system Organization



Operations Performed on Directory

- ▶ Search for a file
- ▶ Create a file
- ▶ Delete a file
- ▶ List a directory
- ▶ Rename a file
- ▶ Traverse the file system




Organize the Directory (Logically) to Obtain

- ▶ Efficiency – locating a file quickly
- ▶ Naming – convenient to users
 - Two users can have same name for different files
 - The same file can have several different names
- ▶ Grouping – logical grouping of files by properties, (e.g., all Java programs, all games, ...)



Single-Level Directory

- ▶ A single directory for all users




Naming problem

Grouping problem



Two-Level Directory


- ▶ Separate directory for each user



- ✓ Path name
- ✓ Can have the same file name for different user
- ✓ Efficient searching
- ✓ No grouping capability



Tree-Structured Directories



Tree-Structured Directories (Cont)

- ▶ Efficient searching
- ▶ Grouping Capability
- ▶ Current directory (working directory)
 - `cd /spell/mail/prog`
 - `type list`



Tree-Structured Directories (Cont)

- ▶ **Absolute** or **relative** path name
- ▶ Creating a new file is done in current directory
- ▶ Delete a file


`rm <file-name>`

- ▶ Creating a new subdirectory is done in current directory

`mkdir <dir-name>`

Example: if in current directory `/mail`

`mkdir count`




Deleting “mail” $\Rightarrow$ deleting the entire subtree rooted by “mail”



Acyclic-Graph Directories

- ▶ Have shared subdirectories and files



Acyclic-Graph Directories (Cont.)

- ▶ Two different names (aliasing)
- ▶ If *dict* deletes *list* $\Rightarrow$ dangling pointer

Solutions:

- Backpointers, so we can delete all pointers
Variable size records a problem
- Backpointers using a daisy chain organization
- Entry-hold-count solution
- ▶ New directory entry type
 - **Link** – another name (pointer) to an existing file
 - **Resolve the link** – follow pointer to locate the file





General Graph Directory (Cont.)

- ▶ How do we guarantee no cycles?
 - Allow only links to file not subdirectories
 - Garbage collection
 - Every time a new link is added use a cycle detection algorithm to determine whether it is OK

