

Opportunity

Virtual Memory 1

Memory references are dynamically translated into physical addresses at run time

A process may be swapped in and out of main memory such that it occupies different regions

A process may be broken up into pieces that do not need to be located contiguously in main memory

Not all pieces of a process need to be loaded in main memory at once during execution

Execution of a Program

Virtual Memory 2

Operating system brings into main memory a few pieces of the program

Resident set - portion of process that is in main memory

An interrupt is generated when an address is needed that is not in main memory – *page fault*

Operating system places the process in a blocking state

Piece of process that contains the logical address is brought into main memory

Operating system issues a disk I/O Read request

Another process is dispatched to run while the disk I/O takes place

An interrupt is issued when disk I/O complete which causes the operating system to place the affected process in the Ready state

Advantages of Breaking up a Process

Virtual Memory 3

More processes may be maintained in main memory

Only load in some of the pieces of each process

With so many processes in main memory, it is very likely at least one process will be in the Ready/Run state at any particular time

A process may be larger than all of main memory

Terminology

Virtual Memory 4

Real memory physical memory, RAM, main memory

Virtual memory secondary storage holding process images

Thrashing phenomenon that a process is spending more time paging than executing

Locality program and data references within a process tend to cluster;
implies that only a few pieces of a process will be needed over a short period of time;
possible to make intelligent guesses about which pieces will be needed in the future;
suggests that virtual memory may work efficiently

Paging

Virtual Memory 5

Each process has its own *page table*

Each page table entry contains the frame number of the corresponding page in main memory

A *resident bit* is needed to indicate whether the page is currently in main memory

Modify bit is needed to indicate if the page has been altered since it was last loaded into main memory

If no change has been made, the page does not have to be written to the disk when it needs to be swapped out

virtual address

Page Number	Offset
-------------	--------

page table entry

r	m	...	Frame Number
---	---	-----	--------------

Computer Science Dept Va Tech August 2007
Operating Systems
©2003-07 McQuain

Address Translation

Virtual Memory 6

Frame # is just concatenated with offset to obtain physical address

If page size is 2^k ...

Page # is leftmost $n = 32 - k$ bits of virtual address...

... so no arithmetic is necessary to extract it...

The diagram illustrates the address translation process across three stages: Program, Paging Mechanism, and Main Memory.

- Program:** A **Virtual Address** is shown, split into **Page #** (leftmost n bits) and **Offset**.
- Paging Mechanism:**
 - The **Page #** is used to look up the **Frame #** in the **Page Table**.
 - The **Page Table** entry contains the **Frame #** (m bits).
 - The **Page #** and **Frame #** are concatenated to form the **Physical Address**.
 - The **Offset** is also part of the **Physical Address**.
- Main Memory:** The **Physical Address** is used to locate the **Page Frame** in memory. The **Offset** is used to find the specific byte within that frame.

Computer Science Dept Va Tech August 2007
Operating Systems
©2003-07 McQuain

Cost of Page Indexing

Virtual Memory 7

The entire page table may take up too much main memory

- each process may typically be allowed a virtual memory space of 2 GB or more
- given 4 KB pages, and a 4GB virtual space, there could be 2^{20} page table entries for each process
- each page table entry would occupy, say, 4 bytes of space, so the page table would occupy 4 MB of memory (per process), or 2^{10} frames/pages

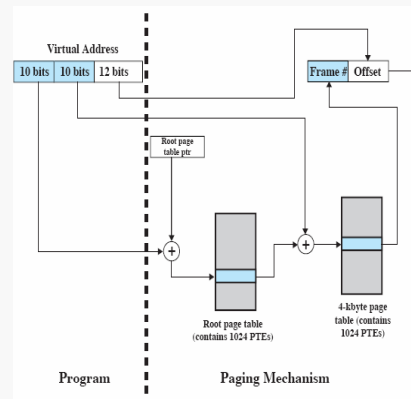
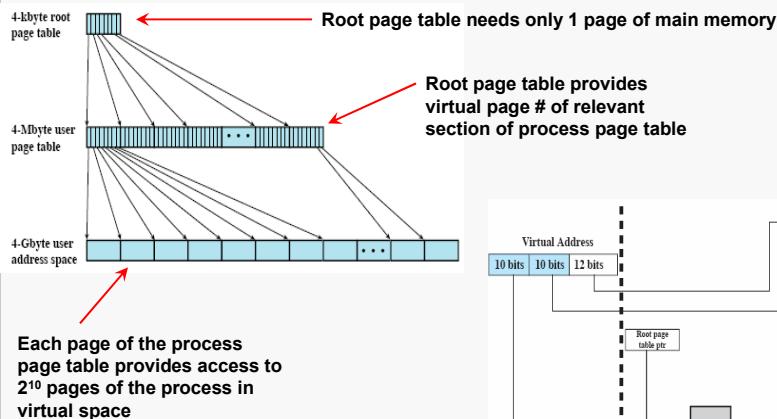
Therefore, page tables are also stored in virtual memory and loaded into main memory as needed.

The 2^{20} page table entries can be efficiently indexed using a 2-level structure:

- the 2^{10} pages of the page table can be indexed via a root page table with 2^{10} entries, needing only 4 KB (one page) of main memory
- lock the root page table in main memory...

Two-Level Scheme for 32-bit Address

Virtual Memory 8



The entire page table may take up too much main memory
Page tables are also stored in virtual memory
When a process is running, part of its page table is in main memory

Inverted page table

- page number portion of a virtual address is mapped into a hash value
- hash value points to inverted page table
- fixed proportion of real memory is required for the tables regardless of the number of processes
- used on PowerPC, UltraSPARC, and IA-64 architecture

Each virtual memory reference can cause two physical memory accesses

- one to fetch the page table
- one to fetch the data

To overcome this problem a high-speed cache is set up for page table entries

- called a Translation Lookaside Buffer (*TLB*)
- contains page table entries that have been most recently used

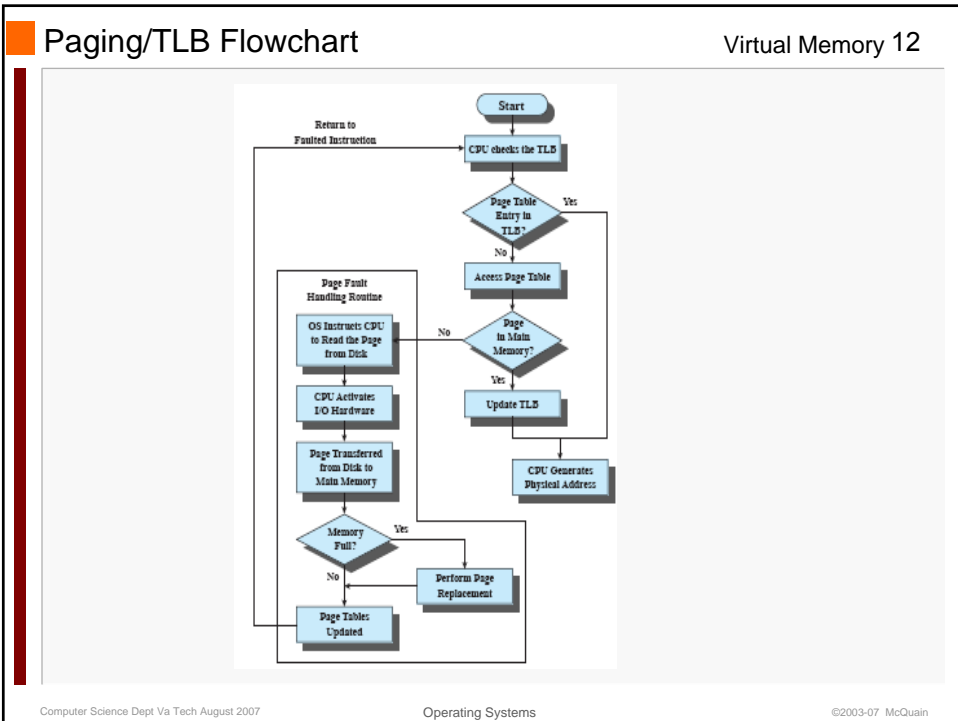
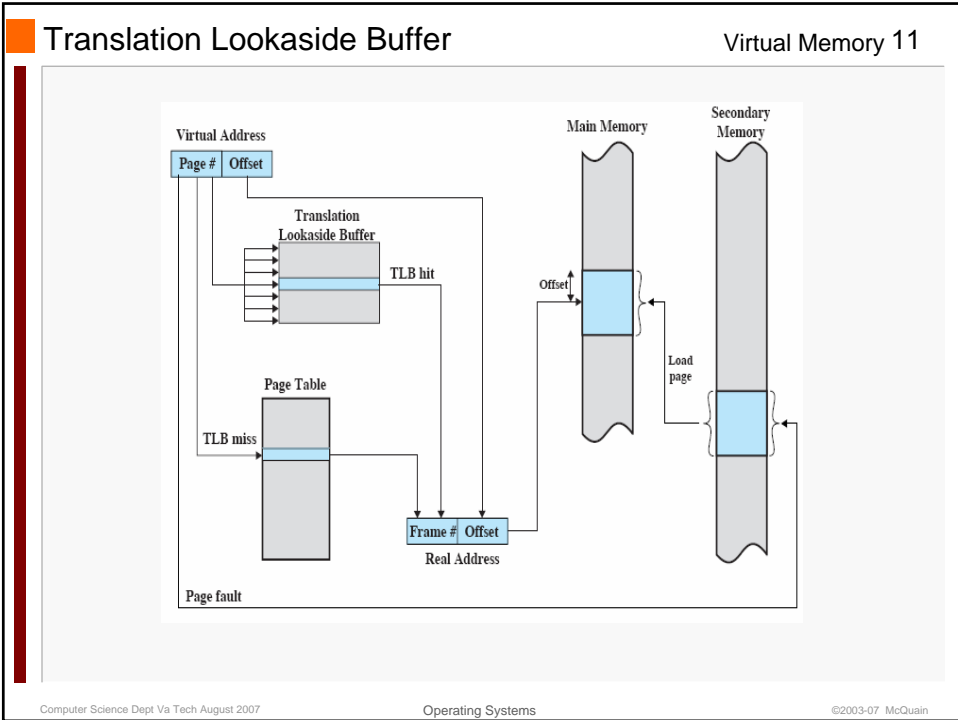
Given a virtual address, processor examines the TLB

If page table entry is present (*TLB hit*):

- the frame number is retrieved and the real address is formed

If page table entry is not found in the TLB (*TLB miss*):

- the page number is used to index the process page table
- if page is already in main memory, proceed
- if not, trigger a page fault
- update TLB to index the new page



Direct Mapping

Page # is index of corresponding entry in the page table, so the relevant table entry can be found in $\Theta(1)$ time via the page table.

The diagram illustrates the direct mapping process. At the top, a box labeled 'Virtual Address' contains 'Page #' with the value 5 and 'Offset' with the value 502. An arrow points from the 'Page #' value to a 'Page Table' represented as a vertical stack of 10 cells. The 5th cell from the top is highlighted in blue and contains the value 37. An arrow points from this cell to a box at the bottom right labeled 'Real Address', which contains 'Frame #' with the value 37 and 'Offset' with the value 502.

Computer Science Dept Va Tech August 2007
Operating Systems
©2003-07 McQuain

Associative Mapping

Unfortunately, the TLB does not contain all of the entries that would be in the page table.

So, direct mapping won't work in the TLB.

But, we still need efficient lookup.

The TLB would support fully associative lookup... in simplest terms, every location in the TLB can be compared to the desired value at once, so lookup will still be $\Theta(1)$.

The diagram illustrates associative mapping. At the top, a box labeled 'Virtual Address' contains 'Page #' with the value 5 and 'Offset' with the value 502. An arrow points from the 'Page #' value to a 'Translation Lookaside Buffer' (TLB) represented as a vertical stack of 10 entries. Each entry has a 'Page #' and 'PT Entries' column. The entries are: (19, empty), (511, empty), (37, empty), (27, empty), (14, empty), (1, empty), (211, empty), (5, 37), (90, empty), and three empty entries below. The entry with Page # 5 and PT Entry 37 is highlighted in blue. An arrow points from this entry to a box at the bottom right labeled 'Real Address', which contains 'Frame #' with the value 37 and 'Offset' with the value 502.

Computer Science Dept Va Tech August 2007
Operating Systems
©2003-07 McQuain

