07 Memory

07.03 Memory hierarchy

- Reference locality
- Caching
- Virtual memory

Virtual memory

- Uses secondary storage to extend the address space beyond physical memory
- The CPU produces *Virtual addresses* that need to be translated into *Physical addresses*
- *Relocation*: a program can run in any location in physical memory
- *Protection*: a process can access only the memory blocks allocated to it

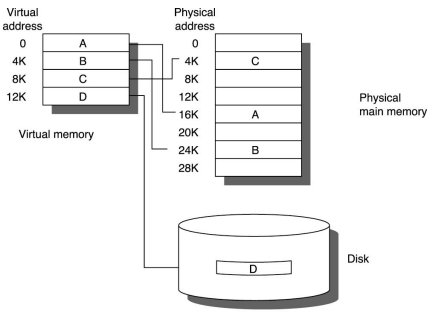

© 2003 Elsevier Science (USA). All rights reserved.

Virtual memory *vs* caches

Caching	Virtual Memory
Cache	Main Memory
Main Memory	Disk
Block	Segment / Page
Cache miss	Page fault
Replacement mainly controlled by HW	Raplecement mainly controlled by SW

Parameter	L1 Cache	Virtual Memory
Block (page) size	16-128 bytes	4 - 64 Kbyte
Hit time	1-3 Tclk	50-150 Tclk
Miss penalty	8-150 Tclk	10^6-10^7 Tclk
(access time)	6-130 Tclk	8E5-8E6 Tclk
(transfer time)	2-20 Tclk	2E5-2E6 Tclk
Miss rate	0.1-10%	0.00001-0.001%
Mapping	25-45 bit to 14-20	32-64 bit to 25-45

Segments *vs* pages

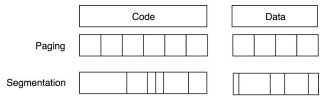

© 2003 Elsevier Science (USA). All rights reserved.

	Page	Segment
Block size	Constant	Variable
Words per address	One	Two (segment and offset)
Visibility	Invisible to application programmer	May be visible to application programmer
Replacement	Trivial (all blocks are the same size)	Hard (must find contiguous, variable-size, unused space in main memory)
Fragmentation	Internal (unused portion of a page)	External (unused space between segments)
Disk traffic	Efficient (adjust page size to balance access time and transfer time)	Inefficient (small segments may transfer only a few bytes)

Virtual Mem: page placement

- Address miss involves access to secondary storage
  - Very long miss penalty imposes to reduce miss rate
- Memory hit of non-cached addresses takes tens of clock cycles
  - Memory access time hides block identification time
- The block placement strategy in main memory is typically fully-associative
  - Replacement policy is usually LRU

Virtual Mem: page identification

- Both paging and segmentation rely on data structures indexed by the page or segment number

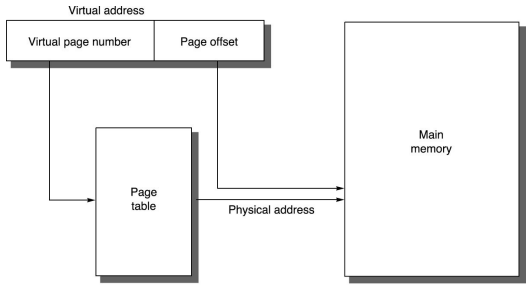

© 2003 Elsevier Science (USA). All rights reserved.

Address translation

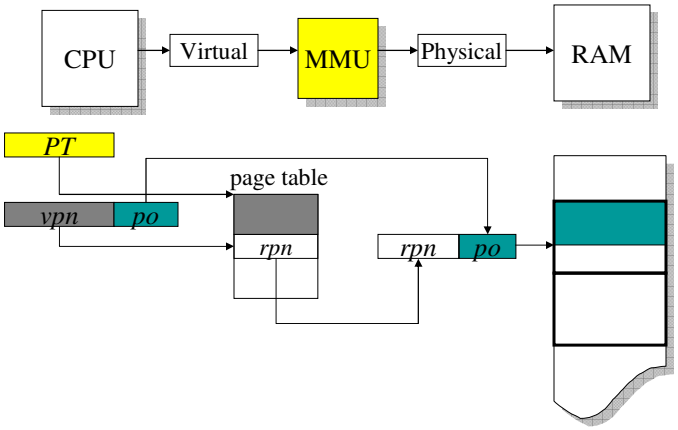

Address translation
(paged page table)

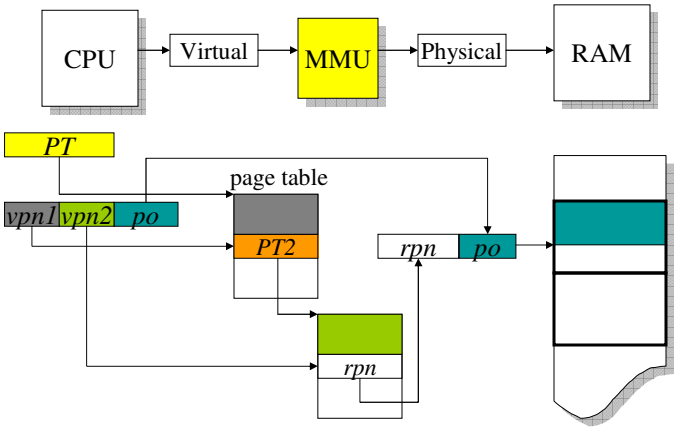

Address translation
(Translation Lookaside Buffer, TLB)

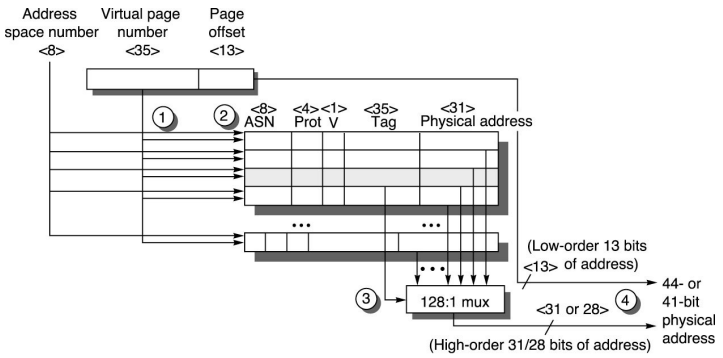

© 2003 Elsevier Science (USA). All rights reserved.

Cache and Virtual memory

Physical cache: address translation needed to access the cache

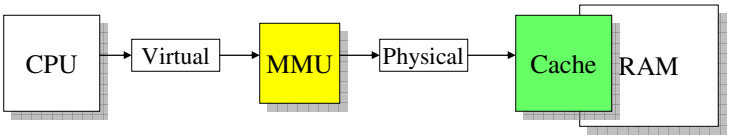

Virtual cache: no address translation needed to access the cache