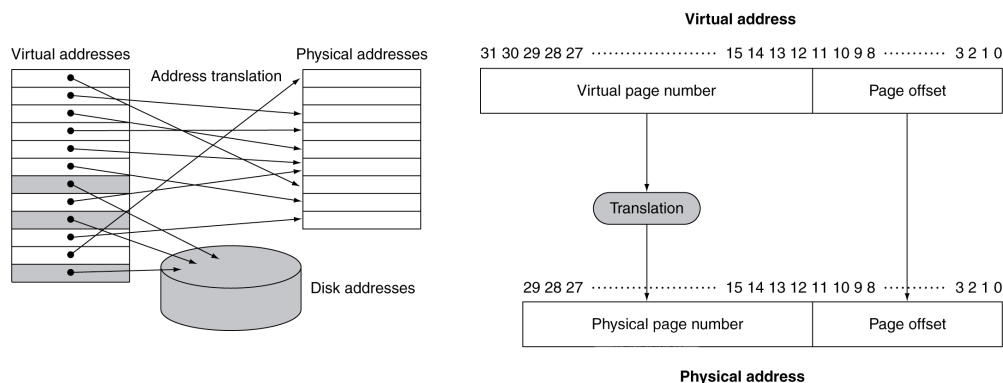


## Virtual Memory

- ❑ Use main memory as a “cache” for secondary (disk) storage
  - Managed jointly by CPU hardware and the operating system (OS)
- ❑ Programs share main memory
  - Each gets a private virtual address space holding its frequently used code and data
  - Protected from other programs
- ❑ CPU and OS translate virtual addresses to physical addresses
  - VM “block” is called a page
  - VM translation “miss” is called a page fault

## Address Translation

- ❑ Fixed-size pages (e.g., 4K)



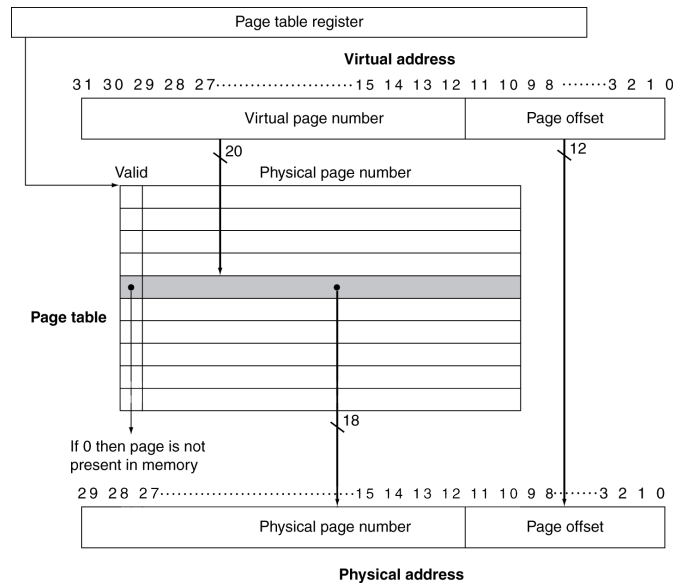
## Page Fault Penalty

- ❑ On page fault, the page must be fetched from disk
  - Takes millions of clock cycles
  - Handled by OS code
- ❑ Try to minimize page fault rate
  - Fully associative placement
  - Smart replacement algorithms

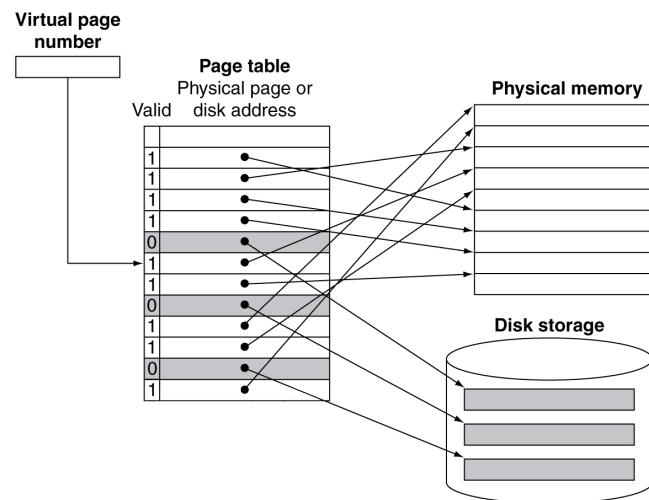
## Page Tables

- ❑ Stores placement information
  - Array of page table entries, indexed by virtual page number
  - Page table register in CPU points to page table in physical memory
- ❑ If page is present in memory
  - PTE stores the physical page number
  - Plus other status bits (referenced, dirty, ...)
- ❑ If page is not present
  - PTE can refer to location in swap space on disk

# Translation Using a Page Table



## Mapping Pages to Storage



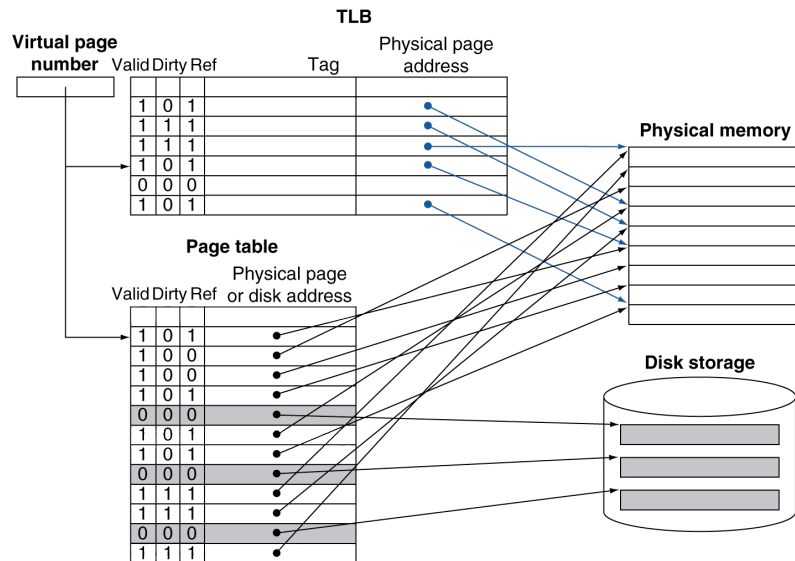
## Replacement and Writes

- ❑ To reduce page fault rate, prefer least-recently used (LRU) replacement
  - Reference bit (aka use bit) in PTE set to 1 on access to page
  - Periodically cleared to 0 by OS
  - A page with reference bit = 0 has not been used recently
- ❑ Disk writes take millions of cycles
  - Block at once, not individual locations
  - Write through is impractical
  - Use write-back
  - Dirty bit in PTE set when page is written

## Fast Translation Using a TLB

- ❑ Address translation would appear to require extra memory references
  - One to access the PTE
  - Then the actual memory access
- ❑ But access to page tables has good locality
  - So use a fast cache of PTEs within the CPU
  - Called a Translation Look-aside Buffer (TLB)
  - Typical: 16–512 PTEs, 0.5–1 cycle for hit, 10–100 cycles for miss, 0.01%–1% miss rate
  - Misses could be handled by hardware or software

## Fast Translation Using a TLB



## TLB Misses

- ❑ If page is in memory
  - Load the PTE from memory and retry
  - Could be handled in hardware
    - Can get complex for more complicated page table structures
  - Or in software
    - Raise a special exception, with optimized handler
- ❑ If page is not in memory (page fault)
  - OS handles fetching the page and updating the page table
  - Then restart the faulting instruction

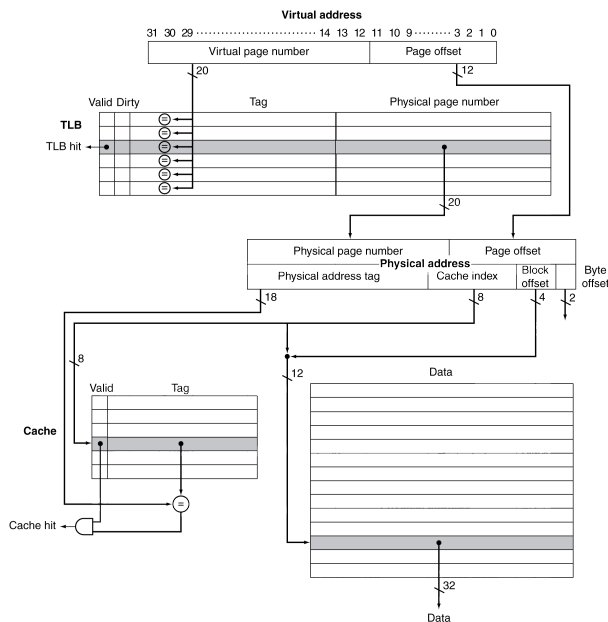
## TLB Miss Handler

- ❑ TLB miss indicates
  - Page present, but PTE not in TLB
  - Page not present
- ❑ Must recognize TLB miss before destination register overwritten
  - Raise exception
- ❑ Handler copies PTE from memory to TLB
  - Then restarts instruction
  - If page not present, page fault will occur

## Page Fault Handler

- ❑ Use faulting virtual address to find PTE
- ❑ Locate page on disk
- ❑ Choose page to replace
  - If dirty, write to disk first
- ❑ Read page into memory and update page table
- ❑ Make process runnable again
  - Restart from faulting instruction

## TLB and Cache Interaction



CS-281 Page 13

Bressoud Spring 2010

- ❑ If cache tag uses physical address
  - Need to translate before cache lookup
- ❑ Alternative: use virtual address tag
  - Complications due to aliasing
    - Different virtual addresses for shared physical address

## Memory Protection

- ❑ Different tasks can share parts of their virtual address spaces
  - But need to protect against errant access
  - Requires OS assistance
- ❑ Hardware support for OS protection
  - Privileged supervisor mode (aka kernel mode)
  - Privileged instructions
  - Page tables and other state information only accessible in supervisor mode
  - System call exception (e.g., syscall in MIPS)

CS-281 Page 14

Bressoud Spring 2010