


# CS 202 Advanced Operating Systems


Virtual Memory

# OS Abstractions



# The CPU-Memory Gap

The gap widens between DRAM, disk, and CPU speeds.



# Locality to the Rescue!




The key to bridging this CPU-Memory gap is a fundamental property of computer programs known as **locality**

# Locality

- **Principle of Locality:** Programs tend to use data and instructions with addresses near or equal to those they have used recently


- **Temporal locality:**

- Recently referenced items are likely to be referenced again in the near future



- **Spatial locality:**

- Items with nearby addresses tend to be referenced close together in time



# Locality Example

```
sum = 0;
for (i = 0; i < n; i++)
    sum += a[i];
return sum;
```


## › Data references

- › Reference array elements in succession (stride-1 reference pattern). Spatial locality
- › Reference variable `sum` each iteration. Temporal locality

## › Instruction references

- › Reference instructions in sequence. Spatial locality
- › Cycle through loop repeatedly. Temporal locality

# An Example Memory Hierarchy




# Memory hierarchy



- ❓ **Cache:** A smaller, faster storage device that acts as a staging area for a subset of the data in a larger, slower device.
- ❓ Fundamental idea of a memory hierarchy:
  - u For each layer, faster, smaller device caches larger, slower device
  - .
- ❓ Why do memory hierarchies work?
  - u Because of locality!
    - › Hit fast memory much more frequently even though its smaller
  - u Thus, the storage at level  $k+1$  can be slower (but larger and cheaper!)
- ❓ **Big Idea:** The memory hierarchy creates a large pool of storage that costs as much as the cheap storage near the bottom, but that serves data to programs at the rate of the fast storage near the top.


# Virtual Addresses



- Many ways to do this translation...
  - Need hardware support and OS management algorithms
- Requirements
  - Need protection – restrict which addresses jobs can use
  - Fast translation – lookups need to be fast
  - Fast change – updating memory hardware on context switch

# Paging

- New Idea: split virtual address space into multiple partitions
  - Each can go anywhere!



Paging solves the external fragmentation problem by using fixed sized units in both physical and virtual memory

But need to keep track of where things are!

# Process Perspective




- ▶ Processes view memory as one contiguous address space from 0 through N
  - ▶ Virtual address space (VAS)
- ▶ In reality, pages are scattered throughout physical storage
- ▶ The mapping is invisible to the program
- ▶ Protection is provided because a program cannot reference memory outside of its VAS
  - ▶ The address “0x1000” maps to different physical addresses in different processes

# Paging




- › Translating addresses
  - › Virtual address has two parts: **virtual page number** and **offset**
  - › Virtual page number (VPN) is an index into a page table
  - › Page table determines page frame number (PFN)
  - › Physical address is PFN::offset
- › Page tables
  - › Map **virtual page number** (VPN) to **page frame number** (PFN)
    - › VPN is the index into the table that determines PFN
  - › One page table entry (PTE) per page in virtual address space
    - › Or, one PTE per VPN

# Page Lookups



# VM as a Tool for Caching

- *Virtual memory* is an array of  $N$  contiguous bytes stored on disk.
- The contents of the array on disk are cached in *physical memory* (*DRAM cache*)
  - These cache blocks are called *pages* (size is  $P = 2^p$  bytes)




# DRAM Cache Organization



- ▶ DRAM cache organization driven by the enormous miss penalty
  - ▶ DRAM is about **10x** slower than SRAM
  - ▶ Disk is about **10,000x** slower than DRAM
- ▶ Consequences
  - ▶ Large page (block) size: typically 4-8 KB, sometimes 4 MB
  - ▶ Fully associative
    - ▶ Any VP can be placed in any PP
    - ▶ Requires a “large” mapping function – different from CPU caches
  - ▶ Highly sophisticated, expensive replacement algorithms
    - ▶ Too complicated and open-ended to be implemented in hardware
  - ▶ Write-back rather than write-through


# Page Tables

- A *page table* is an array of page table entries (PTEs) that maps virtual pages to physical pages.
- Per-process kernel data structure in DRAM




# Page Hit

- *Page hit*: reference to VM word that is in physical memory (DRAM cache hit)




# Page Fault

- **Page fault:** reference to VM word that is not in physical memory (DRAM cache miss)




# Handling Page Fault

- Page miss causes page fault (an exception)




# Handling Page Fault

- Page miss causes page fault (an exception)
- Page fault handler selects a victim to be evicted (here VP 4)




# Handling Page Fault

- Page miss causes page fault (an exception)
- Page fault handler selects a victim to be evicted (here VP 4)



# Handling Page Fault

- › Page miss causes page fault (an exception)
- › Page fault handler selects a victim to be evicted (here VP 4)
- › Offending instruction is restarted: page hit!




# Locality to the Rescue!



- › Virtual memory works because of locality
- › At any point in time, programs tend to access a set of active virtual pages called the *working set*
  - › Programs with better temporal locality will have smaller working sets
- › If (working set size < main memory size)
  - › Good performance for one process after compulsory misses
- › If ( SUM(working set sizes) > main memory size )
  - › *Thrashing*: Performance meltdown where pages are swapped (copied) in and out continuously


# VM as a Tool for Mem Management

- Key idea: each process has its own virtual address space
  - It can view memory as a simple linear array
  - Mapping function scatters addresses through physical memory
    - Well chosen mappings simplify memory allocation and management



# VM as a Tool for Mem Management

- Memory allocation
  - Each virtual page can be mapped to any physical page
  - A virtual page can be stored in different physical pages at different times
- Sharing code and data among processes
  - Map virtual pages to the same physical page (here: PP 6)



# Copy on Write



- › OSes spend a lot of time copying data
  - › System call arguments between user/kernel space
  - › Entire address spaces to implement fork()
- › Use Copy on Write (CoW) to defer large copies as long as possible, hoping to avoid them altogether
  - › Instead of copying pages, create **shared mappings** of parent pages in child virtual address space
  - › Shared pages are protected as read-only in parent and child
    - › Reads happen as usual
    - › Writes generate a protection fault, trap to OS, copy page, change page mapping in client page table, restart write instruction
  - › How does this help fork()?

# Execution of fork()

Parent process's  
page table


|        |
|--------|
| Page 1 |
| Page 2 |
|        |

Physical Memory



|        |
|--------|
| Page 1 |
| Page 2 |
|        |


Child process's  
page table



# fork() with Copy on Write

When either process modifies Page 1,  
page fault handler allocates new page  
and updates PTE in child process


Parent process's  
page table




Physical Memory



Protection bits set to prevent either  
process from writing to any page



Child process's  
page table




# Simplifying Linking and Loading

## › Linking

- › Each program has similar virtual address space  
0xc0000000
- › Code, stack, and shared libraries always start at the same address

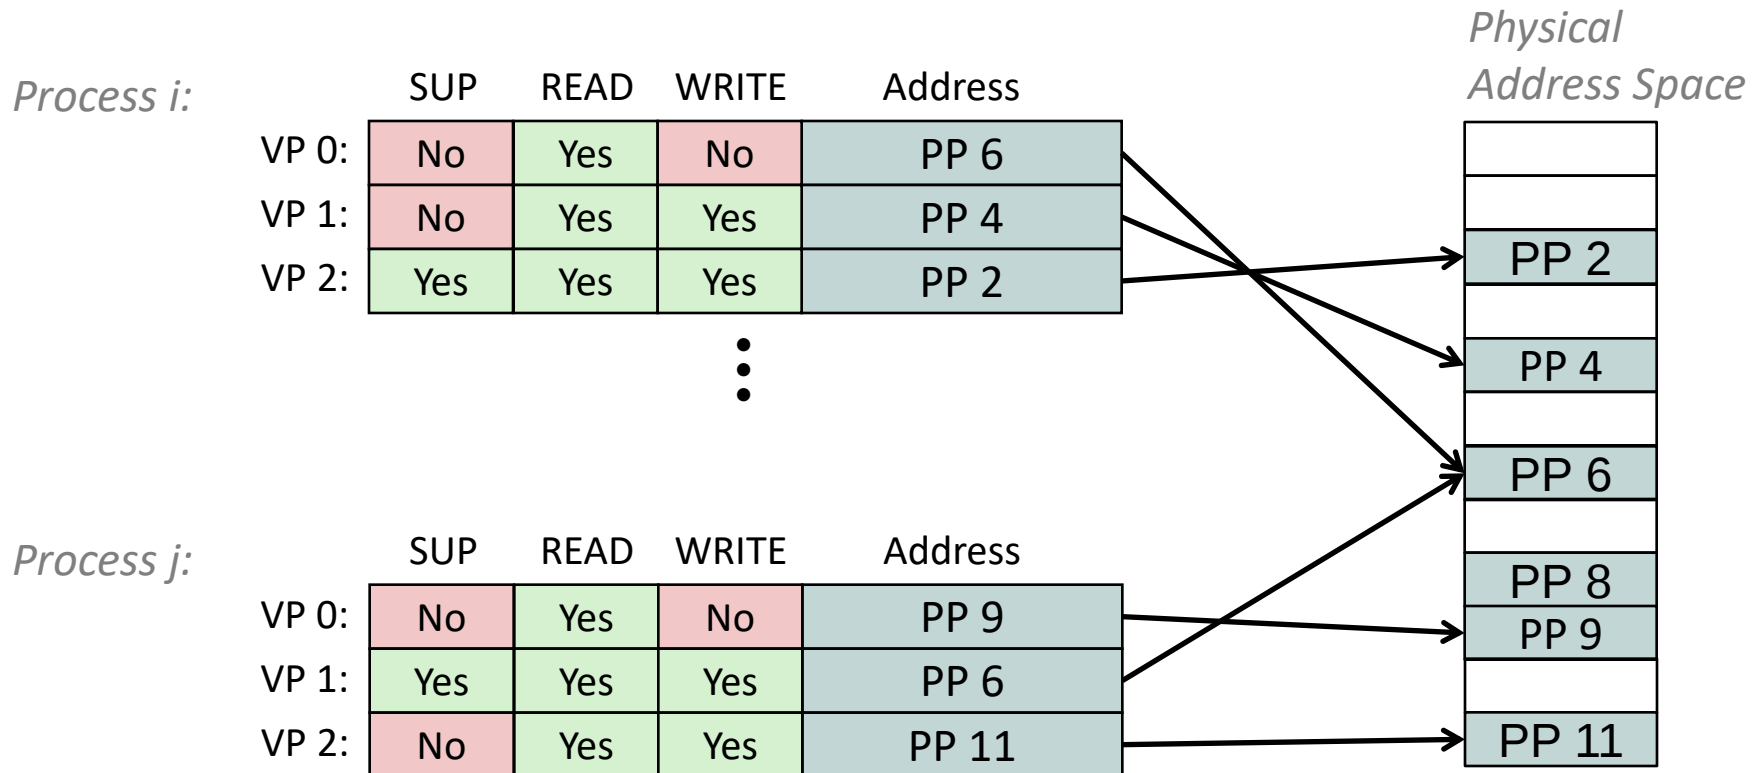
## › Loading

- › **execve()** allocates virtual pages for .text and .data sections  
0x40000000  
= creates PTEs marked as invalid
- › The **.text** and **.data** sections are copied, page by page, on demand by the virtual memory system




# VM as a Tool for Mem Protection

- ▶ Extend PTEs with permission bits
- ▶ Page fault handler checks these before remapping
  - ▶ If violated, send process SIGSEGV (segmentation fault)



# Integrating VM and Cache



*VA: virtual address, PA: physical address, PTE: page table entry, PTEA = PTE address*