Virtual Memory I

Instructor:

Mark Wyse

Teaching Assistants:

Kyrie Dowling Catherine Guevara Ian Hsiao Jim Limprasert Armin Magness Allie Pfleger Cosmo Wang Ronald Widjaja FIGURING OUT WHY MY HOME SERVER KEEPS RUNNING OUT OF SWAP SPACE AND CRASHING:



PLUGGING IT INTO A LIGHT TIMER SO IT REBOOTS EVERY 24 HOURS:



5 MINUTES

WHY EVERYTHING I HAVE IS BROKEN

https://xkcd.com/1495/

Roadmap



Virtual Memory (VM*)

- Overview and motivation
- * VM as a tool for caching
- Address translation
- VM as a tool for memory management
- VM as a tool for memory protection

Warning: Virtual memory is pretty complex, but crucial for understanding how processes work and for debugging performance

*Not to be confused with "Virtual Machine" which is a whole other thing.

Hop

Jate

cere

OxFF·····F

0x00.....0

Memory as we know it so far... is virtual!

- Programs refer to virtual memory addresses
 - movq (%rdi),%rax
 - Conceptually memory is just a very large array of bytes
 - System provides private address space to each process
- Allocation: Compiler and run-time system
 - Where different program objects should be stored
 - All allocation within single virtual address space
- ✤ But...
 - We *probably* don't have 2^w bytes of physical memory
 - We *certainly* don't have 2^w bytes of physical memory <u>for every process</u>
 - Processes should not interfere with one another
 - Except in certain cases where they want to share code or data

Problem 1: How Does Everything Fit?



Problem 2: Memory Management

Physical main memory



Problem 3: How To Protect

Physical main memory



Problem 4: How To Share?

Physical main memory



How can we solve these problems?

 "Any problem in computer science can be solved by adding another level of indirection." – David Wheeler, inventor of the subroutine



What if I want to move Thing?

Indirection

- Indirection: The ability to reference something using a name, reference, or container instead of the value itself. A flexible mapping between a name and a thing allows changing the thing without notifying holders of the name.
 - Adds some work (now have to look up 2 things instead of 1)
 - But don't have to track all uses of name/address (single source!)
- Examples:
 - **Phone system:** cell phone number portability
 - Domain Name Service (DNS): translation from name to IP address
 - Call centers: route calls to available operators, etc.
 - Dynamic Host Configuration Protocol (DHCP): local network address assignment

Indirection in Virtual Memory



- Each process gets its own private virtual address space
- Solves the previous problems!

Address Spaces

- * Virtual address space: Set of $N = 2^n$ virtual addr
 - {0, 1, 2, 3, ..., N-1}
- * Physical address space: Set of $M = 2^m$ physical addr
 - {0, 1, 2, 3, ..., M-1}
- Every byte in main memory has:
 - one physical address (PA)
 - zero, one, or more virtual addresses (VAs)

Polling Questions

On a 64-bit machine currently running 8 processes, how much virtual memory is there?

True or False: A 32-bit machine with 8 GiB of RAM installed would never use all of it (in theory).

Mapping

- A virtual address (VA) can be mapped to either physical memory or disk
 - Unused VAs may not have a mapping
 - VAs from *different* processes may map to same location in memory/disk



Summary

- Virtual memory provides:
 - Ability to use limited memory (RAM) across multiple processes
 - Illusion of contiguous virtual address space for each process
 - Protection and sharing amongst processes