

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Virtual Memory I

CSE 351 Winter 2018

Instructor:
Mark Wyse

Teaching Assistants:
Kevin Bi
Parker DeWilde
Emily Furst
Sarah House
Waylon Huang
Vinny Palaniappan

<http://fabru.com/te/had-chrome-1162082/>

3

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Administrative

- ❖ Homework 4 due tonight
- ❖ Lab 4 due next Wednesday (2/28)

2

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Creating New Processes & Programs

3

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

exit: Ending a process

- ❖ `void exit(int status)`
 - Exits a process
 - Status code: 0 is used for a normal exit, nonzero for abnormal exit

4

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Zombies

- ❖ When a process terminates, it still consumes system resources
 - Various tables maintained by OS
 - Called a "zombie" (a living corpse, half alive and half dead)
- ❖ *Reaping* is performed by parent on terminated child
 - Parent is given exit status information and kernel then deletes zombie child process
- ❖ What if parent doesn't reap?
 - If any parent terminates without reaping a child, then the orphaned child will be reaped by `init` process (`pid == 1`)
 - **Note:** on more recent Linux systems, `init` has been renamed `systemd`
 - In long-running processes (e.g. shells, servers) we need *explicit* reaping

5

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

wait: Synchronizing with Children

- ❖ `int wait(int *child_status)`
 - Suspends current process (*i.e.* the parent) until one of its children terminates
 - Return value is the PID of the child process that terminated
 - *On successful return, the child process is reaped*
 - If `child_status != NULL`, then the `*child_status` value indicates why the child process terminated
 - Special macros for interpreting this status – see `man wait(2)`
- ❖ **Note:** If parent process has multiple children, `wait` will return when *any* of the children terminates
 - `waitpid` can be used to wait on a specific child process

6

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Process Management Summary

- ❖ `fork` makes two copies of the same process (parent & child)
 - Returns different values to the two processes
- ❖ `exec*` replaces current process from file (new program)
 - Two-process program:
 - First `fork()`
 - `if (pid == 0) { /* child code */ } else { /* parent code */ }`
 - Two different programs:
 - First `fork()`
 - `if (pid == 0) { execv(...) } else { /* parent code */ }`
- ❖ `wait` or `waitpid` used to synchronize parent/child execution and to reap child process

7

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Summary

- ❖ Processes
 - At any given time, system has multiple active processes
 - On a one-CPU system, only one can execute at a time, but each process appears to have total control of the processor
 - OS periodically “context switches” between active processes
 - Implemented using *exceptional control flow*
- ❖ Process management
 - `fork`: one call, two returns
 - `execve`: one call, usually no return
 - `wait` or `waitpid`: synchronization
 - `exit`: one call, no return

8

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Roadmap

C:

```
car *c = malloc(sizeof(car));
c->miles = 100;
c->gals = 17;
float mpg = get_mpg(c);
free(c);
```

Java:

```
Car c = new Car();
c.setMiles(100);
c.setGals(17);
float mpg =
    c.getMPG();
```

Memory & data
Integers & floats
x86 assembly
Procedures & stacks
Executables
Arrays & structs
Memory & caches
Processes
Virtual memory
Memory allocation
Java vs. C

Assembly language:

```
get_mpg:
    pushq   %rbp
    movq   %rsp, %rbp
    ...
    popq   %rbp
    ret
```

Machine code:

```
011010000011000
100011010000010000000010
1000100111000010
11000001111101000011111
```

Computer system:

OS: Windows 10, OSX, Ubuntu

9

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

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.

10

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

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^{32} bytes of physical memory
 - We *certainly* don't have 2^{32} bytes of physical memory **for every process**
 - Processes should not interfere with one another
 - Except in certain cases where they want to share code or data

11

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Problem 1: How Does Everything Fit?

64-bit virtual addresses can address several exabytes
(18,446,744,073,709,551.616 bytes)

Physical main memory offers a few gigabytes
(e.g. 8,589,934,592 bytes)

(Not to scale; physical memory would be smaller than the period at the end of this sentence compared to the virtual address space.)

1 virtual address space per process, with many processes...

12

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Problem 2: Memory Management

We have multiple processes:

- Process 1
- Process 2
- Process 3
- ...
- Process n

Each process has...

- stack
- heap
- .text
- .data
- ...

X

Physical main memory

What goes where?

13

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Problem 3: How To Protect

Physical main memory

Process i

Process j

Problem 4: How To Share?

Physical main memory

Process i

Process j

14

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

How can we solve these problems?

- 1) Fitting a huge address space into a tiny physical memory
- 2) Managing the address spaces of multiple processes
- 3) Protecting processes from stepping on each other's memory
- 4) Allowing processes to share common parts of memory

15

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Indirection

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

Without Indirection

With Indirection

What if I want to move Thing?

16

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

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

17

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Indirection in Virtual Memory

Virtual memory

Process 1

Virtual memory

Process n

Physical memory

mapping

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

18

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Address Spaces

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

19

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

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

20

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

A System Using Physical Addressing

- Used in "simple" systems with (usually) just one process:
 - Embedded microcontrollers in devices like cars, elevators, and digital picture frames

21

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

A System Using Virtual Addressing

- Physical addresses are completely invisible to programs
 - Used in all modern desktops, laptops, servers, smartphones...
 - One of the great ideas in computer science

22

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

Why Virtual Memory (VM)?

- Efficient use of limited main memory (RAM)
 - Use RAM as a cache for the parts of a virtual address space
 - Some non-cached parts stored on disk
 - Some (unallocated) non-cached parts stored nowhere
 - Keep only active areas of virtual address space in memory
 - Transfer data back and forth as needed
- Simplifies memory management for programmers
 - Each process "gets" the same full, private linear address space
- Isolates address spaces (protection)
 - One process can't interfere with another's memory
 - They operate in different address spaces
 - User process cannot access privileged information
 - Different sections of address spaces have different permissions

23

UNIVERSITY of WASHINGTON L20: Virtual Memory I CSE351, Winter 2018

VM and the Memory Hierarchy

- Think of virtual memory as array of $N = 2^n$ contiguous bytes
- Pages of virtual memory are usually stored in physical memory, but sometimes spill to disk
 - Pages are another unit of aligned memory (size is $P = 2^p$ bytes)
 - Each virtual page can be stored in any physical page (no fragmentation!)

24

Address Translation: Page Tables

- ❖ CPU-generated address can be split into:
 - n-bit address: Virtual Page Number Page Offset
- Request is Virtual Address (VA), want Physical Address (PA)
- Note that Physical Offset = Virtual Offset (page-aligned)
- ❖ Use lookup table that we call the *page table* (PT)
 - Replace Virtual Page Number (VPN) for Physical Page Number (PPN) to generate Physical Address
 - Index PT using VPN: page table entry (PTE) stores the PPN plus management bits (e.g. Valid, Dirty, access rights)
 - Has an entry for every virtual page – why?

Page Table Diagram

- ❖ Page tables stored in physical memory
 - Too big to fit elsewhere – managed by MMU & OS
- ❖ How many page tables in the system?
 - One per process

Page Table Address Translation

Valid bit = 0: page not in memory (page fault)

In most cases, the MMU can perform this translation without software assistance

Page Hit

- ❖ **Page hit:** VM reference is in physical memory

Example: Page size = 4 KB

Virtual Addr: Physical Addr:

VPN: PPN:

Page Fault

- ❖ **Page fault:** VM reference is NOT in physical memory

Example: Page size = 4 KB

Provide a virtual address request (in hex) that results in this particular page fault:

Virtual Addr:

Page Fault Exception

- ❖ User writes to memory location
- ❖ That portion (page) of user's memory is currently on disk

```

int a[1000];
int main ()
{
    a[500] = 13;
}
    
```

80483b7: c7 05 10 9d 04 08 0d movl \$0xd,0x8049d10

- ❖ Page fault handler must load page into physical memory
- ❖ Returns to faulting instruction: mov is executed again!
 - Successful on second try

Handling a Page Fault

- Page miss causes page fault (an exception)

The diagram shows a virtual address pointing to a null entry in the Page Table (DRAM). The Page Table (DRAM) has entries for PTE 0, PTE 3, and PTE 7. PTE 0 and PTE 3 have valid entries pointing to Physical memory (DRAM) pages VP 1, VP 2, VP 7, and VP 4. PTE 7 has a null entry. The Physical memory (DRAM) contains pages VP 1, VP 2, VP 7, and VP 4. The Virtual memory (DRAM/disk) contains pages VP 1, VP 2, VP 3, VP 4, VP 6, and VP 7. A red arrow points from the null entry in PTE 0 to VP 4 in the Physical memory (DRAM).

Handling a Page Fault

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

The diagram shows the same state as slide 37, but with a red arrow pointing from the null entry in PTE 0 to VP 4 in the Physical memory (DRAM), indicating it is the selected victim for eviction.

Handling a Page Fault

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

The diagram shows the same state as slide 37, but with a green arrow pointing from the null entry in PTE 0 to VP 3 in the Physical memory (DRAM), indicating that VP 3 is being loaded into the physical memory.

Handling a 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!

The diagram shows the final state after the page fault is handled. The Page Table (DRAM) now has a valid entry for PTE 0 pointing to VP 3 in the Physical memory (DRAM). The Physical memory (DRAM) contains pages VP 1, VP 2, VP 7, and VP 3. The Virtual memory (DRAM/disk) contains pages VP 1, VP 2, VP 3, VP 4, VP 6, and VP 7.

Peer Instruction Question

- How many bits wide are the following fields?
 - 16 KB pages
 - 48-bit virtual addresses
 - 16 GB physical memory

	VPN	PPN
(A)	34	24
(B)	32	18
(C)	30	20
(D)	34	20

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
- Indirection via address mapping by page tables
 - Part of memory management unit and stored in memory
 - Use virtual page number as index into lookup table that holds physical page number, disk address, or NULL (unallocated page)
 - On page fault, throw exception and move page from swap space (disk) to main memory