Processes II & Virtual Memory I

CSE 351 Summer 2024

Instructor: Ellis Haker

Teaching Assistants:

Naama Amiel Micah Chang Shananda Dokka Nikolas McNamee Jiawei Huang



Administrivia

- Today
 - HW20 due (11:59pm)
 - Lab4 due (11:59pm)
- Friday, 8/9
 - RD23 due (1pm)
 - HW21 due (11:59pm)
- Monday, 8/12
 - RD24 due (1pm)
 - HW22 due (11:59pm)
 - Quiz 3 out (11:59pm)

Lecture Topics

- Processes and context switching
 - Creating new processes
 - fork() and exec*()
 - Ending a process
 - exit(), wait(), waitpid()
 - Zombies
- Virtual Memory (VM)
 - Overview and motivation
 - \circ $\,$ VM as a tool for caching
 - Address translation
 - VM as a tool for memory management
 - VM as a tool for memory protection

fork Example forker returns dill's PID to parent, O to dill

```
void fork1() {
    int x = 1;
    pid_t fork_ret = fork();
    if (fork_ret == 0)
        printf("Child has x = %d\n", ++x);
    else
        printf("Parent has x = %d\n", --x);
        printf("Bye from process %d with x = %d\n", getpid(), x);
}
```

- Both parent and child start/continue execution after fork
- Child gets a copy of parent's data both processes start with x = 1
 - Subsequent changes to x are **independent**
- Shared open files stdout is the same for both
- Can't predict execution order of parent and child up to the OS!

Modeling fork with Process Graphs

- A process graph is a useful tool for capturing the partial ordering of statements in a concurrent program
 - Each vertex is the execution of a statement
 - \circ a \rightarrow b means a happens before b
 - Edges can be labeled with current value of variables
 - printf vertices can be labeled with output
 - Each graph begins with a vertex with no in-edges
- Any <u>topological sort</u> of the graph corresponds to a feasible total ordering Total ordering of vertices where all edges point from left to right i.e. any sort where vertices come after any vertex w/ a peth to it



Polling Question



Fork-Exec

- fork() creates a copy of the current process
- exec*() replaces the current process' code and address space with the code for a different program
 - Whole family of exec calls see exec(3) and execve(2)

```
void fork_exec(char* path, char* argv[]) {
    pid_t fork_ret = fork();
    if (fork_ret != 0) {
        printf("Parent: created a child %d\n", fork_ret);
    } else {
        printf("Child: about to exec a new program\n");
        execv(path, argv);
    }
    printf("This line printed by parent only!\n");
}
```

exec-ing a Program

Very high-level diagram of what happens when you run the command "ls" in a Linux shell

• This is the loading part of CALL!

Parent
Stack
Неар
Data
Code: /usr/bin/bash

exec-ing a Program (pt 2)

Very high-level diagram of what happens when you run the command "ls" in a Linux shell

• This is the loading part of CALL!



exec-ing a Program (pt 3)

Very high-level diagram of what happens when you run the command "ls" in a Linux shell

Child

This is the loading part of CALL! Ο



Lecture Topics

- Processes and context switching
 - Creating new processes
 - fork() and exec*()
 - Ending a process
 - exit(), wait(), waitpid()
 - Zombies
- Virtual Memory (VM)
 - Overview and motivation
 - \circ $\,$ VM as a tool for caching
 - Address translation
 - VM as a tool for memory management
 - VM as a tool for memory protection

exit: Exiting a Process

- void exit(int status)
 - Explicitly exits a process
 - Status code: 0 = normal exit, nonzero = abnormal exit
- The return statement from main() also exits a process
 - The return value is the status code

- Terminated processes still take up system resources
 - Data structures maintained by the OS
 - A process can't clean up all of its own resources when it exits, so whose responsibility is it?

Zombies

- A terminated process that is still consuming resources is called a zombie
- Parent needs to reap its zombie children (i.e. clean up its resources)
 - Parent is given exit status information, then transfers control to the OS to delete zombie process
- What if the parent exits before reaping the child?
 - Orphaned child is reaped by init process (process 1)
 - Note: on recent Linux systems, init has been renamed to systemd



wait: Synchronizing with Children

- int wait(int* child_status)
 - Suspends the current process until one of its children terminates
 - Reaps that child, then returns its PID
 - If child_status != NULL, then the *child_status value indicates why the child process terminated
 - If NULL, that means the status was ignored
 - 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

wait Example exit printf void fork_wait() { int child_status; if (fork() == 0) { **"HP"** "ст" "Bye" printf("HC: hello from child\n"); exit(0); لانهل posewajt printf fork printf printf else { printf("HP: hello from parent\n"); Feasible output: Infeasible output: ourent wait(&child_status); HC HP printf("CT: child has terminated\n"); HP CT < printf("Bye\n"); CT Bye } Bye

"HC"

wait Example 2: Zomples		Zombie child is still there	
<pre>void fork7() { if (fork() == 0) { /* Child */ printf("Terminating Child, PID = %d\n", getpid()); exit(0); } else {</pre>	linux> ./forks 7 & [1] 6639 Running Parent, PID Terminating Child, linux> ps PID TTY	there = 6639 PID = 6640 TIME CMD	
<pre>/* Parent */ printf("Running Parent, PID = %d\n", getpid()); while (1); /* Infinite loop */ } } </pre>	6585 ttyp9 00:0 6639 ttyp9 00:0 6640 ttyp9 00:0 6641 ttyp9 00:0 linux> kill 6639 [1] Terminated	0:00 tcsh 0:03 forks 0:00 forks <defunct> 0:00 ps</defunct>	
Need to kill parent for init to reap the child	PID TTY 6585 ttyp9 00:0 6642 ttyp9 00:0	TIME CMD 0:00 tcsh 0:00 ps	

wait Example 3: Non-terminating Child



Lecture Topics

- Processes and context switching
 - Creating new processes
 - fork() and exec*()
 - Ending a process
 - exit(), wait(), waitpid()
 - Zombies
- 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

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

0xF...F Programs refer to **virtual** memory addresses System provides private addresses for each process Ο • Allocation: compiler and run-time system Where different program objects should be stored Ο All allocation within single virtual address space 0 But... We *probably* don't have 2^w bytes of physical memory Ο We *definitely* don't have 2^w bytes of physical memory *for every* Ο process Processes should not interfere with each other Ο 0x0...0 Except for specific cases where they want to share code or data

Problem 1: How does everything fit?

64-bit <u>virtual</u> addresses can address 18 exabytes (18,446,744,073,709,551,616 bytes)



<u>Physical</u> 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.)

Problem 2: Memory Management



Problem 3: How to protect data?



What if two running programs both use the same address in their code?

We want to make sure processes don't access the same physical memory locations.

Problem 4: How to share data?



... Except sometimes we *do* want them to share memory!

- Inter-process communication
- Shared code
- etc.

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 How to me more "thing" to "new thing"? Without indirection: p1 thing p2 change all new 3 pointers thing p3 p1 thing With indirection: p2 only change l pointer : new thing p3

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!)
- <u>Examples</u>:
 - Phone system: cell phone number portability
 - Domain Name Service (DNS): translation from name to IP address
 - Call centers: route calls to available operators, etc.

Indirection in Virtual Memory



- Each process gets its own private address space
 - Translates to some location in physical memory
- Solves previous problems!

Mapping

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



Address Spaces

- Virtual Address Space: Set of $N = 2^n$ virtual addresses
 - **{0, 1, …,** *N***-1}**
 - Corresponds to word size (so in x86-64, n = 64)
- **Physical Address Space:** Set of $M = 2^m$ physical addresses
 - {0, 1, ..., *M*-1}
 - Address length *m* depends on hardware
- Every byte in main memory has:
 - **One** physical address (PA)
 - Zero, one, or more virtual addresses (VAs)

Review Questions

1.On a 64-bit machine currently running 8 processes, how much virtual memory is there? Z^{4} B for P^{2} X B for Z^{4} B

2. True or False: A <u>32-bit</u> machine with <u>8</u> GiB of RAM installed would never use all of it (in theory). = z³² B vistral memory performs conty enough For Z processes! Fotse

VM and the Memory Hierarchy

- Think of memory (virtual or physical) as an array of bytes, now split into pages
 - Pages aligned (size is P = 2p bytes), similar to cache blocks
 - Each virtual page can be stored in any physical page (no fragmentation!)
- Pages of virtual memory are usually stored in physical memory, but spill to disk when we run out of space
 - Kind of like a cache!



Memory Hierarchy: Core 2 Duo



Virtual Memory Design Consequences

- Large page size: typically 4-8 KiB or 2-4 MiB
 - Can be up to 1 GiB (for "Big Data" apps on big computers)
 - Much larger than cache blocks
- Fully associative
 - Any virtual page can be placed in any physical page
- Highly sophisticated, expensive replacement algorithms in OS
 - Too complicated and open-ended to be implemented in hardware
- Write-back rather than write-through
 - *Really* don't want to write to disk every time we modify memory
 - Some things may never end up on disk (e.g., stack for short-lived process)

Why does VM work on RAM/disk?

- Avoids disk accesses because of *locality*
 - Same reason that L1 / L2 / L3 caches work
- The set of virtual pages that a program is "actively" accessing at any point in time is called its **working set**
 - If (working set of one process \leq physical memory):
 - Good performance for one process (after compulsory misses)
 - If (working sets of all processes > physical memory):
 - Thrashing: Performance meltdown where pages are swapped between memory and disk continuously
- This is why your computer can feel faster when you add RAM

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)
- **exit** or return from main to end a process
- wait or waitpid used to synchronize parent/child execution and to reap child
- 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