CS 422/522 Design & Implementation of Operating Systems

Lecture 4: Memory Management & The Programming Interface

Zhong Shao Dept. of Computer Science Yale University

This lecture

To support multiprogramming, we need "Protection"

- Kernel vs. user mode
- What is an address space?
- How to implement it?

Physical memory No protection

Limited size Sharing visible to programs

Abstraction: virtual memory

Each program isolated from all others and from the OS

Illusion of "infinite" memory

Transparent --- can't tell if memory is shared















n(x)	physical men	nory for strlen(x)
108, r2	x: 108	a b c \0
oc+8, r31		
60		
	Main: 4240	store 1108, r2
	4244	store pc+8, r31
e (r2), r3 🛛	4248	jump 360
	424c	
31)		
	strlen: 4360	loadbyte (r2), r3
		,
	4420	jump (r31)
	n (x) 08, r2 c+8, r31 50 : (r2), r3 31)	n (x) 108, r2 c+8, r31 50 (r2), r3 (r2), r3 (r3), r3 (r2), r3 (r3), r3 (r), r3 (r

Paging

Motivations

- both branch bounds and segmentation still require fancy memory management (e.g., first fit, best fit, re-shuffling to coalesce free fragments if no single free space is big enough for a new segment)
- can we find something simple and easy
- Solution
 - allocate physical memory in terms of fixed size chunks of memory, or **pages**.
 - Simpler because it allows use of a bitmap
 00111110000001100 ---- each bit represents one page of physical memory
 1 means allocated, 0 means unallocated



















Summary (cont'd)

- Goal: multiprogramming with protection + illusion of "infinite" memory
- Today's lecture so far:
 - HW-based approach for protection: dual mode operation + address space
 - Address translation: virtual address -> physical address
- Future topics
 - how to make address translation faster? use cache (TLB)
 - demand paged virtual memory
- The rest of today's lecture:
 - The programming interface



Abstraction: process & file system

Problem

- Multiple CPU cores, many I/O devices and lots of interrupts
- Users feel they have machine to themselves

Answer

- Decompose hard problems into simple ones
- Deal with one at a time
- Process is such a unit (reflecting something dynamic)
- File system is another high-level abstraction (for "data")

Future

- How processes differ from threads? What is a process really?
- Generalizing "processes" to "containers" & "virtual machines"

Simplest process

- Sequential execution
 - No concurrency inside a process
 - Everything happens sequentially
 - Some coordination may be required

Process state

- Registers
- Main memory
- I/O devices
 - * File system
 - * Communication ports





Process control block (PCB)

- Process management info
 - State
 - * Ready: ready to run
 * Punning: currently ru
 - Running: currently running
 Blocked: waiting for resources
 - Registers, EFLAGS, and other CPU state
 - Stack, code and data segment
 - Parents, etc
- Memory management info
 - Segments, page table, stats, etc

I/O and file management Communication ports, directories, file descriptors, etc.

How OS takes care of processes
 Resource allocation and process state transition

Primitives of processes

- Creation and termination
 - Exec, Fork, Wait, Kill
- ♦ Signals
 - Action, Return, Handler
- Operations
 - Block, Yield
- Synchronization
 - We will talk about this later

Make a process

Creation

- Load code and data into memory
- Create an empty call stack
- Initialize state to same as after a process switch
- Make the process ready to run
- Clone
 - Stop current process and save state
 - Make copy of current code, data, stack and OS state
 - Make the process ready to run

UNIX process management

- UNIX fork system call to create a copy of the current process, and start it running
 - No arguments!
- UNIX exec system call to change the program being run by the current process
- UNIX wait system call to wait for a process to finish
- UNIX signal system call to send a notification to another process





Implementing UNIX fork & exec

- Steps to implement UNIX fork
 - Create and initialize the process control block (PCB) in the kernel
 - Create a new address space
 - Initialize the address space with a copy of the entire contents of the address space of the parent
 - Inherit the execution context of the parent (e.g., any open files)
 - Inform the scheduler that the new process is ready to run
- Steps to implement UNIX exec
 - Load the program into the current address space
 - Copy arguments into memory in the address space
 - Initialize the hardware context to start execution at ``start''







Scheduling policies

Scheduling issues

- fairness: don't starve process
- prioritize: more important first
- deadlines: must do by time 'x' (car brakes)
- optimization: some schedules » faster than others
- No universal policy:
 - many variables, can't maximize them all
 - conflicting goals
 - * more important jobs vs starving others
 - * I want my job to run first, you want yours.
- Given some policy, how to get control ?



- Traps: events generated by current process
 - system calls
 - errors (illegal instructions)
 - page faults
- Interrupts: events external to the process
 - I/O interrupt
 - timer interrupt (every 100 milliseconds or so)
- Process perspective:
 - explicit: process yields processor to another
 - implicit: causes an expensive blocking event, gets switched

UNIX I/O --- a key innovation ("files")

- Uniformity
 - All operations on all files, devices use the same set of system calls: open, close, read, write
- Open before use
 - Open returns a handle (file descriptor) for use in later calls on the file
- Byte-oriented
- Kernel-buffered reads/writes
- Explicit close
 - To garbage collect the open file descriptor
- ◆ Pipes (for interprocess communication → a kernel buffer with two file descriptors, one for reading, one for writing)

