


COS 318: Operating Systems


Processes and Threads

Jaswinder Pal Singh
Computer Science Department
Princeton University


(<http://www.cs.princeton.edu/courses/cos318/>)



Next Few Lectures




- ◆ Processing: Concurrency and Sharing
 - Processes and threads
 - Synchronization
 - CPU scheduling
 - Deadlock




2

Today's Topics




- ◆ Concurrency
- ◆ Processes
- ◆ Threads




3

Concurrency, Processes and Threads



- ◆ Concurrency
 - Many things going on in an operating system
 - Application process execution, interrupts, background tasks, maintenance
 - CPU is shared, as are I/O devices
 - Human beings are not very good at keep track of this and programming it monolithically
 - Processes (and threads) are abstraction to bridge this gap
- ◆ Concurrency via Processes
 - Decompose complex problems into simple ones
 - Make each simple one a process
 - Processes run 'concurrently' but each process feels like it has its own CPU
- ◆ Example: gcc (via "gcc -pipe -v") launches the following
 - /usr/libexec/cpp | /usr/libexec/cc1 | /usr/libexec/as | /usr/libexec/elf/ld
 - Each instance of cpp, cc1, as and ld running is a process



4

Threads


- ◆ A process has an address space and resources
- ◆ Thread
 - A sequential execution stream within a process (also called lightweight process)
 - Separately schedulable: OS or runtime can run or suspend at any time
 - A process can have one or more threads (loci of execution)
 - Threads in a process share the same address space
- ◆ Can have concurrency across processes, and/or across threads within a process.
 - We will talk in terms of processes first (i.e. every process has only one thread)
 - We will talk about threads after that



5

Process Concurrency

- ◆ Virtualization
 - Processes interleaved on CPU
- ◆ I/O concurrency
 - I/O for P1 overlapped with CPU for P2
 - Each runs almost as fast as if it has its own computer
 - Reduce total completion time
- ◆ CPU parallelism
 - Multiple CPUs (such as SMP)
 - Processes running in parallel
 - Speedup



6

Parallelism

- ◆ Parallelism is common in real life
 - A single sales person sells \$1M annually
 - Hire 100 sales people to generate \$100M revenue
- ◆ Speedup
 - Ideal speedup is factor of N
 - Reality: bottlenecks + coordination overhead reduce speedup
- ◆ Questions
 - Can you speed up by working with a partner?
 - Can you speed up by working with 20 partners?
 - Can you get super-linear (more than a factor of N) speedup?



7


Concurrency in Computing

- ◆ Parallel programs
 - To achieve better performance
- ◆ Servers (expressing logically concurrent tasks)
 - Multiple connections handled simultaneously
- ◆ Programs with user interfaces
 - To achieve user responsiveness while doing computation
- ◆ Network and disk bound programs
 - To hide network/disk latency



The CPU Illusion

- ◆ Each process thinks it owns the CPU
 - yet on a uni-processor all processes share the same physical CPU
 - How does this work?



- ◆ Two key pieces:


- PCB --- process control block, one per process, holds execution state
- dispatching loop:


```
while(1)
  interrupt
  save state
  get next process
  load state, jump to it
```



The Abstraction

- ◆ Infinite number of processors
- ◆ Processes/threads execute with variable speed
 - Programs must be designed to work with any schedule




Programmer vs. processor view

Programmer's View	Possible Execution #1	Possible Execution #2	Possible Execution #3
.	.	.	.
.	.	.	.
x = x + 1;	x = x + 1;	x = x + 1;	x = x + 1;
y = y + x;	y = y + x;	y = y + x;
z = x + 5y;	z = x + 5y;	Thread is suspended.
.	.	Other thread(s) run.	Thread is suspended.
.	.	Thread is resumed.	Other thread(s) run.
.	Thread is resumed.
.	.	y = y + x;
.	.	z = x + 5y;	z = x + 5y;




Possible executions


One Execution



Another Execution (3 CPUs)



Another Execution




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
 - ...



15

Program and Process



16

Process vs. Program

- ◆ Process > program
 - Program is just the code; just part of process state
 - Example: many users can run the same program
- ◆ Process < program
 - A program can invoke more than one process
 - Example: Fork off processes
 - Many processes can be running the same program



17

Process Control Block (PCB)

- ◆ Process management info
 - Identification
 - State
 - Ready: ready to run.
 - Running: currently running.
 - Blocked: waiting for resources
 - Registers, EFLAGS, EIP, 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.
- ◆ Resource allocation and accounting information



18

Process Control Block

Process management	Memory management	File management
Registers Program counter Program status word Stack pointer Process state Priority Scheduling parameters Process ID Parent process Process group Signals Time when process started CPU time used Children's CPU time Time of next alarm	Pointer to text segment Pointer to data segment Pointer to stack segment	Root directory Working directory File descriptors User ID Group ID

Potential fields of a PCB



API for Process Management

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



20

Create A Process

- ◆ Creation
 - Load code and data into memory
 - Create an empty call stack
 - Initialize state
 - Make the process ready to run
- ◆ Cloning a process
 - Save state of current process
 - Make copy of current code, data, stack and OS state
 - Make the process ready to run



21

Unix Example

- ◆ Methods to make processes:
 - fork clones a process
 - exec overlays the current process

```
pid = fork();  
if (pid == 0)  
    /* child process */  
    exec("foo"); /* does not return */  
Else  
    /* parent */  
    wait(pid); /* wait for child to die */
```



22

Fork and Exec in Unix

```

    pid = fork();
    if (pid == 0)
        exec("foo");
    else
        wait(pid);
  
```

```

    foo:
    Main()
    {
    ...
    }
  
```

```

    pid = fork();
    if (pid == 0)
        exec("foo");
    else
        wait(pid);
  
```

```

    pid = fork();
    if (pid == 0)
        exec("foo");
    else
        wait(pid);
  
```

Wait

23

More on Fork

- Parent process has a PCB and an address space
- Create and initialize PCB
- Create an address space
- Copy the content of the parent address space to the new address space
- Inherit the execution context of the parent (e.g. open files)
- Inform scheduler that new process is ready

24

Process Context Switch

- Save a context (everything that a process may damage)
 - All registers (general purpose and floating point)
 - All co-processor state
 - Save all memory to disk?
 - What about cache and TLB?
- Start a context
 - Does the reverse
- Challenge
 - OS code must save state without changing any state
 - E.g. how should OS run without touching any registers?
 - CISC machines have a special instruction to save and restore all registers on stack
 - RISC: reserve registers for kernel or have way to carefully save one and then continue

26

(Reduced) Process State Transition

```

    graph TD
      Create --> Ready
      Ready -- Scheduler Dispatch --> Running
      Running -- Wait for resource --> Blocked
      Blocked -- Resource becomes available --> Ready
      Running -- Terminate --> End
  
```

Running: executing now
 Ready: waiting for CPU
 Blocked: waiting for I/O or lock


27

Threads

- ◆ Thread
 - A sequential execution stream within a process (also called lightweight process)
 - Separately schedulable: OS or runtime can run or suspend at any time
 - A process may have one or more threads (loci of execution)
 - Threads in a process share the same address space
- ◆ Thread concurrency
 - Easier to program overlapping I/O and CPU with threads than with signals
 - Human being likes to do several things at a time
 - A server (e.g. file server) serves multiple requests
 - Multiple CPUs sharing the same memory




Threads (cont'd)



Every thread has its own stack



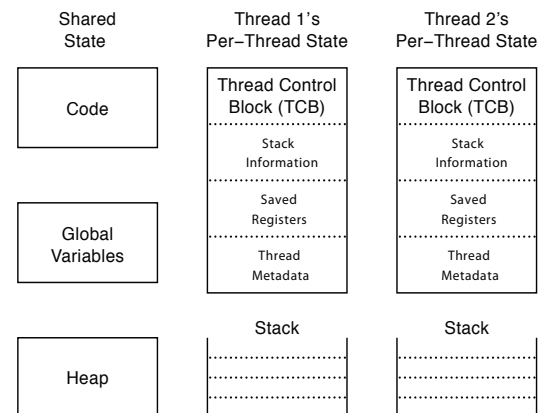
Thread usage



A word processor with three threads



Thread data structures



Thread Control Block (TCB)

- State
 - Ready: ready to run
 - Running: currently running
 - Blocked: waiting for resources
- Registers
- Status (EFLAGS)
- Program counter (EIP)
- Stack
- Code



32

Threads (cont'd)

Per process items

Address space
Global variables
Open files
Child processes
Pending alarms
Signals and signal handlers
Accounting information

Per thread items

Program counter
Registers
Stack
State

- ◆ Items shared by all threads in a process
- ◆ Items private to each thread



Typical Thread API


- ◆ Creation
 - Fork, Join
- ◆ Mutual exclusion
 - Acquire (lock), Release (unlock)
- ◆ Condition variables
 - Wait, Signal, Broadcast
- ◆ Alert
 - Alert, AlertWait, TestAlert



34

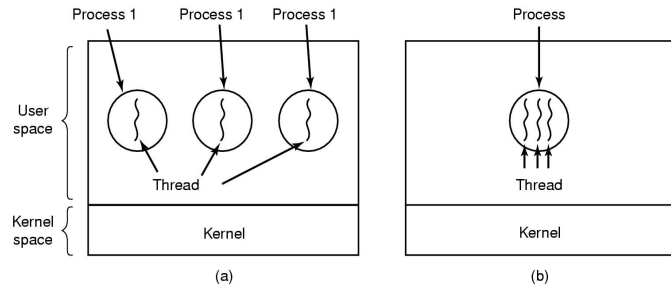
Revisit Process

- ◆ Process
 - Threads
 - Address space
 - Environment for the threads to run on OS (open files, etc)
- ◆ Simplest process has 1 thread



35

Threads and Processes



- (a) Three processes each with one thread
- (b) One process with three threads

- ◆ Process = thread + address space + OS env (open files, etc.)
- ◆ Thread encapsulates concurrency; address space encapsulates protection



38

Thread Context Switch

- ◆ Save a context (everything that a thread may damage)
 - All registers (general purpose and floating point)
 - All co-processor state
 - Need to save stack?
 - What about cache and TLB?
- ◆ Start a context
 - Does the reverse
- ◆ May trigger a process context switch



37

Procedure Call

- ◆ Caller or callee save some context (same stack)
- ◆ Caller saved example:

```

save active caller registers
call foo
    
```

```

foo() {
    do stuff
}
    
```

```

restore caller regs
    
```

Red arrows indicate the flow of control and context saving/restoration between the caller and the callee.



38

Threads vs. Procedures

- ◆ Threads may resume out of order
 - Cannot use LIFO stack to save state
 - Each thread has its own stack
- ◆ Threads switch less often
 - Do not partition registers
 - Each thread "has" its own CPU
- ◆ Threads can be asynchronous
 - Procedure call can use compiler to save state synchronously
 - Threads can run asynchronously
- ◆ Multiple threads
 - Multiple threads can run on multiple CPUs in parallel
 - Procedure calls are sequential



39

Process vs. Threads

- ◆ Address space
 - Processes do not usually share memory (address space)
 - Process context switch page table and other memory mechanisms
 - Threads in a process share the entire address space
- ◆ Privileges
 - Processes have their own privileges (file accesses, e.g.)
 - Threads in a process share all privileges
- ◆ Question
 - Do you really want to share the “entire” address space?



40

Real Operating Systems

- ◆ One or many address spaces
- ◆ One or many threads per address space

	1 address space	Many address spaces
1 thread per address space	MSDOS Macintosh	Traditional Unix
Many threads per address spaces	Embedded OS, Pilot	VMS, Mach (OS-X), OS/2, Windows NT/XP/Vista/7, Solaris, HP-UX, Linux



41

Summary

- ◆ Concurrency
 - CPU and I/O
 - Among applications
 - Within an application
- ◆ Processes
 - Abstraction for application concurrency
- ◆ Threads
 - Abstraction for concurrency within an application



42