

Processes & Threads

CS 475

Concurrent Programs

- ❑ Process = Address space + one thread of control
- ❑ Concurrent program = **multiple threads of control**
 - Multiple single-threaded processes
 - Multi-threaded process

CS 475

2

Concurrent Systems

- ❑ Essential aspects of any concurrent system
 - **Execution context** - state of a concurrent entity
 - Processes: process context
 - Threads: thread context
 - **Scheduling** - deciding which context will run next
 - Processes: Operating System scheduler
 - Threads: Library thread scheduler (Pthreads), Java runtime
 - **Synchronization** - mechanisms that enable execution contexts to coordinate their use of shared resources
 - Semaphores, locks, monitors, condition variables
 - Provided at both operating system and library/language level

Distributed Software Systems

3

Processes

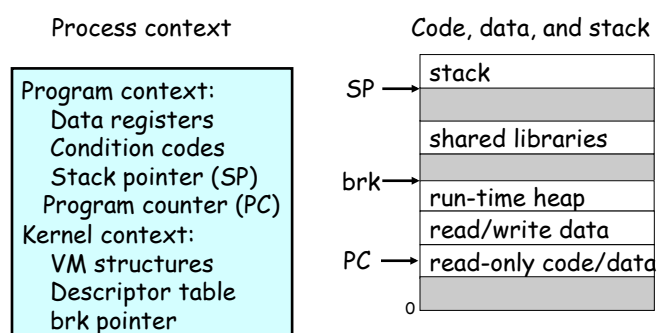
- ❑ Def: A *process* is an instance of a running program.
 - One of the most profound ideas in computer science.
 - Not the same as "program" or "processor"
- ❑ Process provides each program with two key abstractions:
 - Logical control flow
 - Each program seems to have exclusive use of the CPU.
 - Private address space
 - Each program seems to have exclusive use of main memory.
- ❑ How are these illusions maintained?
 - Process executions interleaved (multitasking)
 - Address spaces managed by virtual memory system

CS 475

4

Traditional View of a Process

- Process = process context + code, data, and stack



CS 475

5

Threads: Motivation

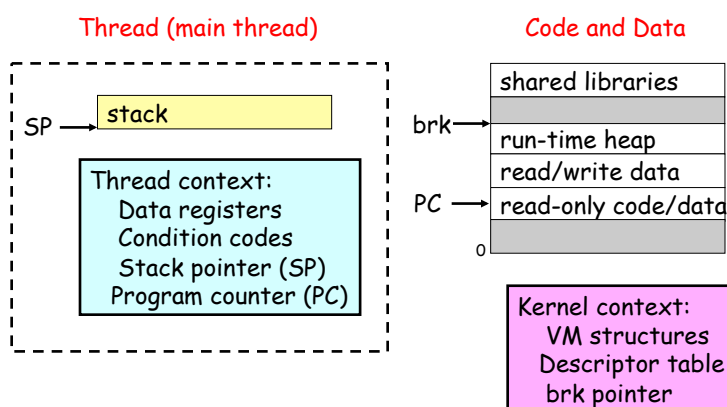
- Traditional processes created and managed by the OS kernel
- Process creation expensive - fork system call in UNIX
- Context switching expensive
- Cooperating processes - no need for memory protection (separate address spaces)

CS 475

6

Alternate View of a Process

- Process = thread + code, data, and kernel context

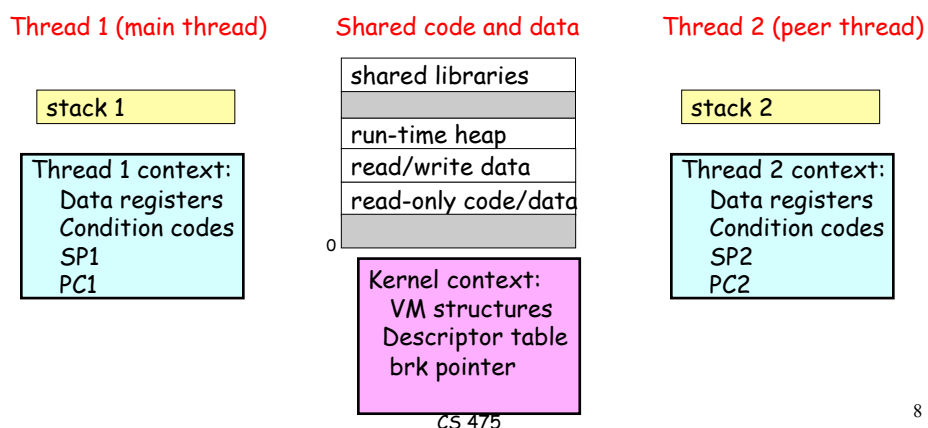


CS 475

7

A Process With Multiple Threads

- Multiple threads can be associated with a process
 - Each thread has its own logical control flow (sequence of PC values)
 - Each thread shares the same code, data, and kernel context
 - Each thread has its own thread id (TID)



CS 475

8

Threads

- ❑ Execute in same address space
 - separate execution stack, share access to code and (global) data
- ❑ Smaller creation and context-switch time
- ❑ Can exploit fine-grain concurrency

CS 475

9

Creating processes

- ❑ UNIX
 - **fork** system call
 - Used in conjunction with **exec** system call

CS 475

10

fork: Creating new processes

❑ `int fork(void)`

- creates a new process (child process) that is identical to the calling process (parent process)
- returns 0 to the child process
- returns child's `pid` to the parent process

```
if (fork() == 0) {
    printf("hello from child\n");
} else {
    printf("hello from parent\n");
}
```

Fork is interesting
(and often confusing)
because it is called
once but returns *twice*

Creating and Using threads

❑ **Pthreads Multi-threading Library**

- API for
- `pthread_create`, `pthread_join`,
`pthread_self`, `pthread_exit`,
`pthread_detach`

❑ **Java**

- provides a `Runnable` interface and a `Thread` class as part of standard Java libraries
 - users program threads by implementing the `Runnable` interface or extending the `Thread` class

Concurrent Systems

- ❑ Essential aspects of any concurrent system
 - **Execution context** - state of a concurrent entity
 - Processes: process context
 - Threads: thread context
 - **Scheduling** - deciding which context will run next
 - Processes: Operating System scheduler
 - Threads: Library thread scheduler (Pthreads), Java runtime
 - **Synchronization** - mechanisms that enable execution contexts to coordinate their use of shared resources
 - Semaphores, locks, monitors, condition variables
 - Provided at both operating system and library/language level

Distributed Software Systems

13

Road Map

- ❑ Next two lectures: Processes & Signals in UNIX
 - Repetition of material discussed in CS 367
 - Assignment 1 (Shell Lab)
- ❑ Thread creation and management in Java and Pthreads (one lecture)
- ❑ Process & Thread synchronization mechanisms (two - three lectures)

CS 475

14