



Parallel & Concurrent Programming: Processes & Threads

Emery Berger
CMPSCI 691W - Spring 2006



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

Outline

- Processes
- Threads
- Basic synchronization
- Bake-off




UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

2

Processes vs. Threads...

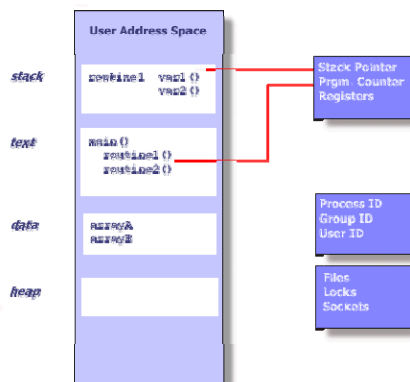
- Both useful for parallel programming & concurrency
 - Hide latency
 - Maximize CPU utilization
 - Handle multiple, asynchronous events
- **But:** different programming styles, performance characteristics, and more




UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

Processes

- **Process:** execution context (PC, registers) + address space, files, etc.
- Basic unit of execution




The diagram illustrates the components of a process. On the left, a vertical stack represents the **User Address Space**, divided into four sections: *stack*, *text*, *data*, and *heap*. The *stack* section contains variables `var1()` and `var2()`. The *text* section contains code `main()`, `routine1()`, and `routine2()`. The *data* section contains `arrayA` and `arrayB`. The *heap* section is empty. On the right, three boxes represent system resources: **Stack Pointer**, **Prpm Counter**, and **Registers** (connected to the stack section); **Process ID**, **Group ID**, and **User ID** (connected to the text section); and **Files**, **Locks**, and **Sockets** (connected to the data section).



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

Process API

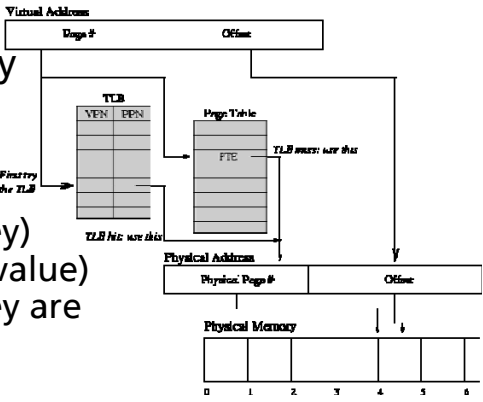
- **UNIX:**
 - `fork()` – create copy of current process
 - Different return value
 - **Copy-on-write**
 - `exec()` – replace process w/ executable
- **Windows:**
 - `CreateProcess (...)`
 - 10 arguments



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

Translation Lookaside Buffer

- TLB: fast, fully associative memory
 - Stores page numbers (key) and frame (value) in which they are stored
- Copy-on-write: **protect pages, copy on first write**



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

Processes Example

```
#include <unistd.h>
#include <sys/wait.h>
#include <stdio.h>
main() {
    int parentID = getpid(); /* ID of this process */
    char prgname[1024];
    gets(prgname); /* read the name of program we want to start */
    int cid = fork();
    if(cid == 0) { /* I'm the child process */
        execlp(prgname, prgname, 0); /* Load the program */
        /* If the program named prgname can be started, we never get
        to this line, because the child program is replaced by prgname */
        printf("I didn't find program %s\n", prgname);
    } else { /* I'm the parent process */
        sleep (1); /* Give my child time to start. */
        waitpid(cid, 0, 0); /* Wait for my child to terminate. */
        printf("Program %s finished\n", prgname);
    }
}
```




Communication

- Processes:
 - Input = state before `fork()`
 - Output = return value
 - argument to `exit()`
- **But:** how can processes communicate *during* execution?



IPC


- **signals**
 - Send & receive ints
 - Not terribly useful for parallel or concurrent programming
- **pipes**
 - Communication channels – easy & fast
 - Just like UNIX command line



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science 9

Pipe example


```
int main() {
    int pfd[2];
    pipe(pfd);
    if (!fork()) {
        close(1); /* close normal stdout */
        dup(pfd[1]); /* make stdout same as pfd[1] */
        close(pfd[0]); /* we don't need this */
        execlp("ls", "ls", NULL);
    } else {
        close(0); /* close normal stdin */
        dup(pfd[0]); /* make stdin same as pfd[0] */
        close(pfd[1]); /* we don't need this */
        execlp("wc", "wc", "-l", NULL);
    }
}
```



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science 10

IPC, continued

- **sockets**
 - Explicit message passing
 - + Can *distribute* processes anywhere
- **shmem**
 - Best not spoken of...
- **mmap (common hack)**
 - All processes map same file into fixed memory location
 - Objects in region shared across processes
 - Use `flock()` to synchronize



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

11

Threads

- **Processes** - everything in distinct address space
- **Threads** - same address space (& files, sockets, etc.)


User Address Space	
Thread 2	<div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">stack</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">main() var1 var2 var3</div>
Thread 1	<div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">stack</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">main() var1 var2</div>
heap	<div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">main() var1 main() var2 ...</div>
data	<div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">main() var1 main() var2</div>
heap	<div style="border: 1px solid black; padding: 2px; margin-bottom: 2px;">main() var1 main() var2</div>

mutex pointer
 Program Counter
 Registers

mutex pointer
 Program Counter
 Registers

Process ID
 User ID
 Group ID

File
 Locks
 Sockets



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science

12

Threads API

■ UNIX (POSIX):

- `pthread_create()` – start separate *thread* executing function
- `pthread_join()` – wait for thread to complete

■ Windows:

- `CreateThread (...)`
 - only 6 arguments!



Threads example

```
#include <pthread.h>
void * run (void * d) {
    int q = ((int) d);
    int v = 0;
    for (int i = 0; i < q; i++) {
        v = v + expensiveComputation(i);
    }
    return (void *) v;
}
main() {
    pthread_t t1, t2;
    int r1, r2;
    pthread_create (&t1, run, 100);
    pthread_create (&t2, run, 100);
    pthread_wait (&t1, (void *) &r1);
    pthread_wait (&t2, (void *) &r2);
    printf ("r1 = %d, r2 = %d\n", r1, r2);
}
```



Communication

- In threads, everything shared except: *stacks, registers & thread-specific data*
 - Old way:
 - `pthread_setspecific`
 - `pthread_getspecific`
 - New way: `__thread`
 - `static __thread int x;`
 - Easier in Java...
- Updates of shared state must be *synchronized*



Basic synchronization


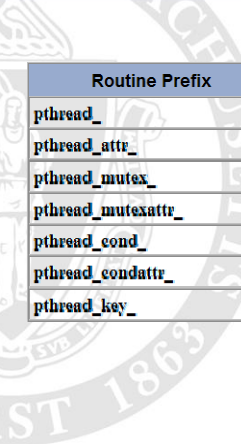
- Mutual exclusion *locks*
 - Only one thread in *critical section*

```
pthread_mutex_lock (&l);  
update data; /* critical section */  
pthread_mutex_unlock (&l);
```



Pthreads API


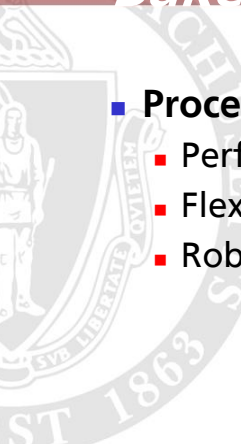
Routine Prefix	Functional Group
pthread_	Threads themselves and miscellaneous subroutines
pthread_attr_	Thread attributes objects
pthread_mutex_	Mutexes
pthread_mutexattr_	Mutex attributes objects.
pthread_cond_	Condition variables
pthread_condattr_	Condition attributes objects
pthread_key_	Thread-specific data keys



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science 17

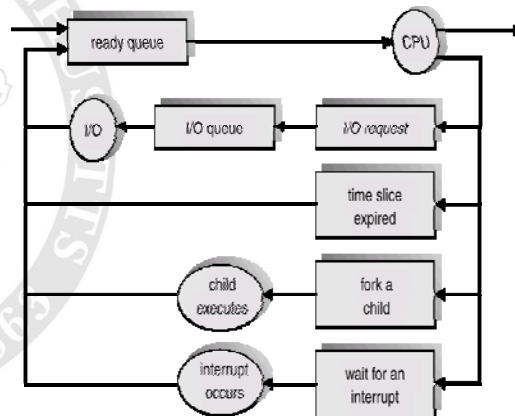
Bake-off

- **Processes or threads?**
 - Performance
 - Flexibility / Ease-of-use
 - Robustness



UNIVERSITY OF MASSACHUSETTS, AMHERST • Department of Computer Science 18

Scheduling



Context Switch Cost

- **Threads** – much cheaper
 - Stash registers, PC (“IP”), stack pointer
 - **Processes:**
 - Same as threads *plus* –
 - Process context
 - **TLB shutdown**
- ⇒ Process switches more expensive, or require long quanta



Flexibility / Ease-of-use

- **Processes** – more flexible
 - + Easy to spawn remotely
 - + Can communicate via sockets = can be distributed across cluster / Internet
 - Requires explicit communication or risky hackery
- **Threads**
 - Communicate through memory – must be on same machine
 - Require *thread-safe* code



Robustness

- **Processes** – far more robust
 - Processes *isolated* from other processes
 - Process dies) no effect
 - Apache 1.x
- **Threads:**
 - If one thread crashes (e.g., derefs NULL), whole process terminates
 - Then there's the stack size problem
 - Apache 2.x...



Next time

- Advanced synchronization

